Fundamenta Informaticae 57 (2003) 281–321
IOS Press

A Modal Characterization of Nash Equilibrium

Paul Harrenstein*
Institute of Information and Computing Sciences
Utrecht University
Utrecht, the Netherlands
paulh@cs.uu.nl

Wiebe van der Hoek
Computer Science Department
The University of Liverpool
Liverpool, United Kingdom
WiebevanderHoek@csc.liv.ac.uk

John-Jules Meyer
Institute of Information and Computing Sciences
Utrecht University
Utrecht, the Netherlands
jj@cs.uu.nl

Cees Witteveen
Faculty of Electrical Engineering, Mathematics and Computer Science
Delft University of Technology
Delft, the Netherlands
C.Witteveen@ewi.tudelft.nl

Abstract. Multi-agent systems comprise entities whose individual decision making behavior may depend on one another's. Game-theory provides apposite concepts to reason in a mathematically precise fashion about such interactive and interdependent situations. This paper concerns a logical analysis of the game-theoretical notions of *Nash equilibrium* and its *subgame perfect* variety as they apply to a particular class of extensive games of perfect information. Extensive games are defined as a special type of labelled graph and we argue that modal languages can be employed in their description. We propose a logic for a multi-modal language and prove its completeness with respect to a class of frames that correspond with a particular class of extensive games. In this multi-modal language (subgame perfect) Nash equilibria can be characterized. Finally, we show how this approach can formally be refined by using Propositional Dynamic Logic (PDL), though we leave completeness as an open question.

Keywords: Modal Logic, Dynamic Logic, Game Theory, Nash Equilibrium

1. Introduction

With the advance of distributed and multi-agent systems there has been an increased interest in the relation between logic and game theory within the field of Artificial Intelligence (*cf.*, *e.g.*, [3], [6] and [26]).

* Paul Harrenstein's research was generously supported by the CABS-project at Delft University of Technology.
Address for correspondence: Institute of Information and Computing Sciences, Utrecht University, Utrecht, the Netherlands

In multi-agent environments, various decision making agents with various degrees of autonomy interact. These individual agents making up a multi-agent system may be designed for the performance of widely divergent and even conflicting tasks. Still, which actions are most conducive to an agent's ends in such situations, may well depend on the decisions of the other agents. The specification and verification of multi-agent systems calls for mathematically precise concepts that facilitate reasoning about such interactive strategic situations. The relevance of game theory to Artificial Intelligence is that it can provide an apposite conceptual framework in this respect.

The theory of games originated in the middle of the 20th Century with the recognition that, to that date, no theory in classical mathematics had dealt with social situations in which each individual tries to maximize a function according to an idiosyncratic principle without having control over all of the variables on which this function depends (*cf.*, [22], p.11). Thus, game theory was developed as the mathematical study of game-like situations in which the eventual outcome depends on the individual choices of various agents, each of which has different preferences over the possible outcomes. In such situations the application of the traditional notions of optimality were thought no longer to suffice and new mathematical concepts were developed to take over their role (*cf.*, *ibid.*, p.39). In this respect, the celebrated *Nash equilibrium* and its *subgame perfect* variety are archetypical in non-cooperative settings. Informally, a collective course of action, or a strategy profile, is said to be a Nash equilibrium if none of the participants has an incentive to deviate unilaterally from that course of action. Whether an agent has an incentive depends on her individual preferences.

One of the guiding ideas of the theory of games is that situations of social interaction can fruitfully be compared with and analyzed as games by distinguishing players, their strategies and their interests. Games have proved to be an especially rewarding metaphor for social environments in which interacting agents are conceived of as players with individual preferences and powers of manipulation. This leaves the question how far the game metaphor goes and how far it should be carried. In order to arrive at a general theory of social interaction, specific and idiosyncratic aspects of games should be abstracted from, in favor of other, more generic features, which should duly be emphasized. Where the dividing line between the general and the specific should be drawn is not an objective matter and may very well depend on one's purposes. Still, it should always be borne in mind that:

> A model structure that is too simple may force us to ignore vital aspects of the real games we want to study. A model structure that is too complicated may hinder our analysis by obscuring fundamental issues. ([19], p.37)

The order in which the players perform their actions in strategic situations has reasonably been argued to be a vital, rather than an obscuring, aspect in this sense. The models of strategic situations provided by games in *extensive form* — or just *extensive games* — are especially designed to account for this type of sequential structure.

One of the important solution concepts that comes along with extensive games is that of *subgame perfect Nash equilibrium*. This ramification of the original notion excludes Nash equilibria that are no longer deemed credible if the sequential structure of games is taken into account. Here, we will give a logical analysis of extensive games and their subgame perfect equilibria. Extensive games are introduced as a special kind of graph. Since modal languages are designed to reason about this kind of relational structure, we employ multi-modal languages and logics to this end. In the second last section we will show how the framework can be refined using Propositional Dynamic Logic (PDL). As such our approach is congenial to Bonanno's in [5], who employed Computational Tree Logic (CTL) to formalize

	left	*right*
top	3 6	4 7
bottom	0 0	1 10

Figure 1.

the concept of backward induction, which is closely related to subgame perfect Nash equilibrium. Also the work of Baltag ([2]) should be mentioned in this context.

Extensive games define a proper subclass of Kripke-frames for the special kind of multi-modal language we consider, as we will argue. A strategy profile for an extensive game corresponds to a subrelation in the frame and as such can be taken as an accessibility relation for a modal operator. Some strategy profiles qualify as a (subgame perfect) Nash equilibrium and others do not. This fact reflects in certain specific structural properties of the modal accessibility relations in the corresponding frame. The result we are after is to characterize these properties by means of a multi-modal formula schema. So, letting G be an extensive game, $\mathfrak{F}_G$ the corresponding frame and s a strategy profile of G, the quest is for a formula schema $\vartheta(s)$ such that:

$$\mathfrak{F}_G \models \vartheta(s) \quad \text{iff} \quad s \text{ is a subgame perfect Nash equilibrium in } G.$$

Any such result would show that subgame perfect Nash equilibrium is a definable property of frames in appropriate multi-modal languages, be it that the frames in question are of a special kind. Accordingly, we suggest an axiomatization of a multi-modal logic of which the semantics is restricted to models based on the class of frames extensive games define. Soundness and completeness, moreover, are proved. Remarkably, the axioms are nothing much out of the ordinary and can be bestowed rather intuitive readings as well. The very austerity of the whole analysis we take as something speaking in its favor. Be that as it may, it accommodates us with a neat modal logic to reason about extensive games and their (subgame perfect) Nash equilibria.

2. Extensive Games with Perfect Information

Generally speaking, putting the egg in the pan first and then the butter does not work quite as well as putting in the butter first and then the egg. The order in which the actions are performed does matter in some cases. In strategic environments this is no different. What is more, the order in which agents can make their choices and moves often makes a *strategic difference* and as such is something the game-theorist had better not ignore entirely.

For an example, consider the strategic situation depicted in Figure 1. Here two players, *Row* and *Col*, choose between rows (*top* or *bottom*) and between columns (*left* or *right*), respectively. The matrix merely summarizes the payoffs to the players for the different possible choices of action. The figure bottom left in each quarter indicates the payoff to *Row*, the figure top right the one awarded to *Col*. The matrix is thought of as specifying no temporal structure whatsoever; it does not even presuppose that the players move simultaneously.

In this particular situation *Row* may be tempted to play *bottom*, the idea being that this would leave *Col* the relatively unfavorable choice between an outcome of one and an outcome of zero. Anticipating that *Col* chooses the former, *Row* may expect a payoff of ten. In order to deter *Row* from taking this course of action, *Col* may threaten to play *left* if *Row* were to play *bottom*; this would result in the worst outcome for both players. This would force *Row* to play *top* which guarantees a better outcome for *Col*. However, if the game structure were such that *Row* is (able) to move first, *Col*'s threats would be rendered void, provided he is not otherwise committed to choose the left column. If in defiance of *Col*'s threats, *Row* were to play bottom anyway, *Col* would be presented with a *fait accompli* and *Col* had better make the best of a bad job and opt for the right colomn after all. Similarly, if *Col* were to move first a threat on his part does not make sense either. In that case, however, he can secure a more favorable outcome by choosing the left column. That would leave *Row* the easy choice between a payoff of six or one of zero. *Row* making the obvious decision would guarantee *Col* a payoff of three, instead of the probable and miserly one he would have gotten had his initial choice been *right* with *Row* quite likely seizing the opportunity and playing *bottom*. This time, however, *Row* can try to achieve a better outcome for both by promising — and committing herself to fulfill this promise — to play *top* if *Col* decides on the right column.

Figure 1 as such leaves unspecified the sequential structure of how the game is played. Above we gave two possible interpretations, an obvious third would be to conceive Figure 1 as a fully-fledged game in strategic form, *e.g.*, by assuming the player to make their choices simultaneously or in ignorance of one another's. Be that as it may, the point of these reflections is that the order in which the players make moves does make a strategic difference. For one thing, the feasibility of making a threat or a promise may depend on it and in our example even the outcome of the game may as well. As such, the sequential structure of game has sensibly been made subject of game-theoretic study.

One way of making the sequential structure of games explicit in a mathematically precise manner is by representing them in their so-called *extensive form*, *i.e.*, as labelled trees. Play commences at the root and each edge indicates a possible course of action for a player. At each node a player is to strike a decision how to act. The two temporal interpretations of Figure 1 can thus be represented as in Figure 2, below. The vectors at the leaves indicate the payoffs to the players, in both cases the first entry being the payoff *Row*, the second to *Col*. Observe that in both cases, there are four strategies for the second player that is to move. In the left picture, each of *Col*'s strategies has to specify whether to play *right* or *left* both if the first player plays *top* and if she plays *bottom*.

Similar concerns may drive the game-theorist to consider models of game-like situations with even more structure. In this manner, strategic reasoning may be argued to depend on epistemic features such as the players' knowledge of the situation they are in or the beliefs they entertain about it or about one another's beliefs, preferences and rationality. Or, he may consider players that randomize over their (pure) strategies. In an effort at keeping our logical analysis as perspicuous as possible, however, we will abstract from these issues and confine our attention to extensive games in pure strategies with perfect information, *i.e.*, the players are assumed to play a (pure) strategy with probability one or zero and they are assumed to be fully informed about the game's structure and the other players' preferences and powers. Moreover, we will assume that only one player can move at a time and that the games will

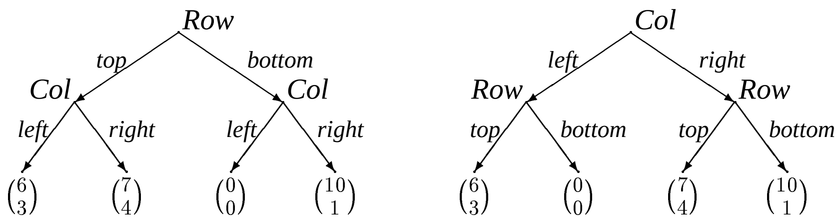

Figure 2.

eventually come to an end after a finite number of moves. With respect to the preferences of the players, we take into account the ordinal structure they determine over the possible outcomes only, which suffices for our purposes. In this paper, concerns as to the intensity of preference as expressible by a specific rational or real number, do not enter the picture. Also disregarding uncertainty on the part of the players as well as well as mixed (or randomized) strategies, our analyses are of a strictly qualitative nature. The following definition makes the notion of *games in extensive form*, or *extensive games*, mathematically precise.

Definition 2.1. (Games in extensive form of perfect information)
A *game in extensive form* G is a tuple $\left(V, R, N, P, \{\rho_i\}_{i \in N}\right)$, where V is a set of vertices (or nodes) and R a relation on V such that (V, R) is a, possibly infinite, directed and irreflexive tree with a finite horizon, *i.e.*, (V, R) contains no infinite branches. The root node of (V, R) is usually denoted by v_0. Furthermore, N is a non-empty but *finite* set of players. The function P assigns to each internal node in V the player in N that has to move at v. Finally, for each i in N, ρ_i is a total pre-order (a reflexive, transitive and connected relation) over the vertices in V, specifying i's preferences. $(v, v') \in \rho_i$ signifies that i values v' at least as high as v. A player i is called *indifferent* if $\rho_i = V \times V$ and *interested*, otherwise. Let Z be the set of leaves of (V, R) and let also, for each player i, V_i denote the subset of vertices in which i is to move, *i.e.*, the set $\{v \in V : P(v) = i\}$.

This definition differs from more conventional ones in that the players' preferences are defined over all vertices rather than over the leaves only. Although for the relevant game-theoretical concepts the preferences over the leaf nodes suffice, we found that defining preferences over all vertices is more convenient for our logical analyses. Note further that the players' preferences over the internal nodes are independent of their preferences over the leaf nodes. In particular, they are not assumed to coincide with the preferences over the internal nodes that backward induction would give rise to.

An extensive game is a labelled tree, the vertices of which represent the possible game positions and the edges (v, v') are possible actions for the player assigned to v. After a player has decided to play along a certain edge and acted accordingly, the game reaches a new game state. In the position then reached, the game either terminates or is again the root node of an extensive game. This idea gives rise to the notion of a *subgame*. Let G be the game $\left(V, R, N, P, \{\rho_i\}_{i \in N}\right)$. For any subtree (V', R') of (V, R) generated by some vertex v, another extensive game is obtained by appropriately restricting the assignment function P and each ρ_i to the vertices in V'. For each vertex v in V we define the subgame, G_v as the tuple $\left(V', R', N, P', \{\rho'_i\}_{i \in N}\right)$, where $V' = \{v' \in V : (v, v') \in R^*\}$, $R' = \left\{(v', v'') \in V' \times V' : (v', v'') \in R\right\}$, $P' = P \upharpoonright V'$ and for each $i \in N$, $\rho'_i = \rho_i \cap (V' \times V')$.

Here R^* denotes the reflexive transitive closure of R.

A *(pure) strategy* for a player in an extensive form is a complete plan for that player to play that game. As such a strategy has to account for a player's choices at all stages of the game in which that player is in control. A strategy even has to prescribe a player's actions in stages of the game it itself precludes from being reached. Intuitively, a *strategy profile* is then a combination of strategies, for each player one. The set of strategy profiles in an extensive game G is denoted by S_G, omitting the subscript where no ambiguity can arise. For our concerns the notion of a strategy profile is more fundamental than that of a strategy. We define a strategy profile s of an extensive game formally as a function mapping each *internal* vertex onto a vertex that succeeds it, *i.e.*, for each v in V, $(v, s(v)) \in R$. For any pair of strategy profiles s and s' and for each subset of players M we have $s^M_{s'}$ denote the strategy profile that is like s except on the vertices assigned to one of the players in M where it takes values from s', *i.e.*, for all internal vertices v we have:

$$s^M_{s'}(v) \quad =_{df.} \quad \begin{cases} s'(v) & \text{if } P(v) \in M, \\ s(v) & \text{otherwise.} \end{cases}$$

Also, $s^i_{s'}$ abbreviates $s^{\{i\}}_{s'}$. A *strategy for a player* i is then the restriction of a strategy profile to the vertices in which i is in control. Accordingly the set of strategies for a player i is defined by $\{s \upharpoonright V_i : \ s \in S\}$.

From each vertex v onwards a strategy profile s generates a path through the game-tree until a leaf node is reached. This path is given by the sequence $v_0, \ldots, v_n$ such that $v_0 = v$, v_n is a leaf and for all $0 \leqslant i < n$, $v_{i+1} = s(v_i)$, *i.e.*, the sequence:

$$v,\ s(v),\ \ldots,\ s^{n-1}(v),$$

where $s^{n-1}(v)$ denotes the $n-1$-fold application of s to v, *e.g.*, $s^3(v) = s(s(s(v)))$. In this manner each strategy profile determines for each vertex a unique leaf node as outcome. With each strategy profile s we accordingly associate an *outcome function*, $\hat{s}$, which maps each vertex on the leaf it has as an outcome if the strategy profile s is followed iteratively. Formally we define for each strategy profile s the outcome function $\hat{s}$ inductively such that for each vertex v in V:

$$\hat{s}(v) \quad =_{df.} \quad \begin{cases} v & \text{if } v \text{ is a leaf,} \\ \hat{s}(s(v)) & \text{otherwise.} \end{cases}$$

On this basis the notion of a strategy profile being a *Nash equilibrium* can be introduced formally. Intuitively, a strategy profile s is a Nash-equilibrium if none of the players benefits from unilaterally deviating from s. Since the strategy profile $s^i_{s'}$ represents the strategy profile that results if i unilaterally deviates from s by playing the strategy prescribed to him in s', we arrive at the following formal definition. For s a strategy profile in an extensive game $G = \left(V, R, N, P, \{\rho_i\}_{i \in N}\right)$:

$$s \text{ is a } \textit{Nash equilibrium} \quad \text{iff} \qquad \text{for all } i \in N, \text{and all } s' \in S_G\colon \left(\hat{s}^i_{s'}(v_0), \hat{s}(v_0)\right) \in \rho_i.$$

As an individual pendant of Nash-equilibrium we have the concept of a *best response for a player* i defined for strategy profiles s as:

$$s \text{ is a } \textit{best response for } i \quad \text{iff} \qquad \text{for all } s' \in S_G\colon \left(\hat{s}^i_{s'}(v_0), \hat{s}(v_0)\right) \in \rho_i.$$

Obviously then, a strategy profile is a Nash-equilibrium if and only if it is a best response for all players.

The notion of a Nash equilibrium entirely focusses on the outcomes the various strategy profiles determine from the root. Different strategy profiles may very well give rise to an identical path from root to leaf node — and as such determine the same outcome — and still differ widely on vertices that are not on this path. The path a Nash equilibrium determines through the game tree is one unilateral deviation from which invariably brings no good to the defector. Yet, it has been argued that if the sequential structure of a game is taken into account, the notion of Nash equilibrium fails to make some important distinctions. Although one could accept — were it only for the sake of argument — the refusal to defect unilaterally from the equilibrium path as the very hallmark of game-theoretical level-headedness, *off* the equilibrium path a Nash equilibrium may strike one as unsatisfactory. Consider once more the extensive game in Figure 2 in which *Row* is to move first. We have already argued that *Row* need not refrain from playing *bottom* even if *Col* were to threaten to choose the left column in that case. This is vindicated by all strategy profiles in which *Row* chooses *bottom* and *Col* subsequently playing *right* being Nash equilibria in this game. Strategies, however, determine choices for the players at all nodes where they are to play. The node that would have been reached had *Row* chosen *top* is no exception. At that node it would be slightly incomprehensible if *Col* were to choose the left column. Still the strategy profile in which *Row* plays *bottom* and *Col* plays *right* if *Row* were to play *bottom* and *left* otherwise, is nevertheless a Nash equilibrium.

As a refinement of Nash equilibrium that however does do justice to the sequential structure of an extensive game, Selten ([29]) proposed the solution concept of a *subgame perfect Nash equilibrium*. Roughly speaking, a strategy profile is a subgame perfect Nash equilibrium in an extensive game G if it is a Nash equilibrium in all subgames of G. In the example above, any strategy profile that would prescribe *Col* to play *left* when *Row* has chosen the top row, would not qualify as a subgame perfect Nash equilibrium. Formally, for G the extensive game $\big(V, R, N, P, \{\rho_i\}_{i\in N}\big)$, define for each strategy profile s:

s is a *subgame perfect Nash equilibrium* iff for all $v \in V, i \in N,$ and $s' \in S_G$: $\big(\hat{s}^i_{s'}(v), \hat{s}(v)\big) \in \rho_i$.

As individual pendant of this concept we also introduce a *subgame perfect best response for a player* i as a strategy profile that is a best response for i in all subgames. Formally, a definition is obtained by omitting the universal quantification over the players in the *definiens* of a subgame perfect Nash-equilibrium.

The following example illustrates the concepts that have been introduced so far.

Example 2.1. Figure 3 gives a graphical representation of a two-player game in extensive form. The preferences of the players over the leaves are represented by the vectors appended to the leaf nodes. The first entry indicates the preferences of Player 1 and the second those of Player 2. The higher the value, the more the outcome is preferred by the player. *E.g.*, (z_5, z_3) is in the preference relation ρ_2, because two is smaller than four. Player 1 has four strategies at her disposal and Player 2 six. Accordingly, there are twenty-four strategy profiles in all, each of which we indicate by a four letter subscript corresponding to the direction the players move at the vertices v_0, v_3, v_1 and v_2, respectively. The choices of player 1 are denoted by capitals, those of player 2 by lower case letters. *E.g.*, the strategy profile s_{RLlr} is the functional relation given by $\big\{(v_0, v_2), (v_3, z_1), (v_1, v_3), (v_2, z_6)\big\}$. Starting from the root v_0, it gives rise

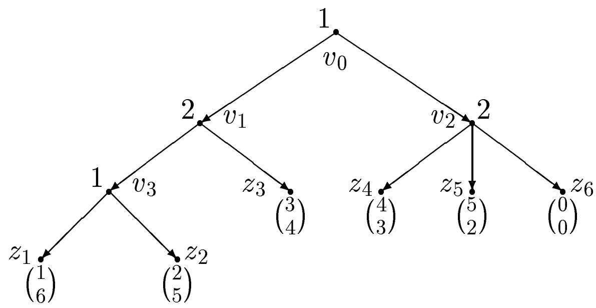

Figure 3.

to the sequence v_0, v_2, z_4 and, accordingly, we have

$$\hat{s}_{RLlr}(v_0) = \hat{s}_{RLlr}(s_{RLlr}(v_0)) = \hat{s}_{RLlr}(v_2) = \hat{s}_{RLlr}(s_{RLlr}(v_2)) = \hat{s}_{RLlr}(z_6) = z_6.$$

This strategy profile, however, fails as a Nash equilibrium. Player 2 could deviate from s_{RLlr} at v_2 and play l there instead. This would make that s_{RLll} would be played, yielding z_4 as outcome and guaranteeing him a payoff of 3 instead of 0. The Nash equilibria of this particular game can be identified with the following relations on the vertices:

$$s_{RLll} = \{(v_0, v_2), (v_3, z_1), (v_1, v_3), (v_2, z_4)\} \qquad s_{RLrl} = \{(v_0, v_2), (v_3, z_1), (v_1, z_3), (v_2, z_4)\}$$
$$s_{RRll} = \{(v_0, v_2), (v_3, z_2), (v_1, v_3), (v_2, z_4)\} \qquad s_{RRrl} = \{(v_0, v_2), (v_3, z_2), (v_1, z_3), (v_2, z_4)\}$$
$$s_{LRlr} = \{(v_0, v_1), (v_3, z_2), (v_1, v_3), (v_2, z_6)\}$$

Of these only s_{RRll} is a subgame perfect Nash equilibrium as well. The strategy profile s_{LRlr}, *e.g.*, is excluded as a subgame perfect equilibrium since it is not a Nash-equilibrium in the subgame that has v_2 as root.

Obviously, every subgame perfect Nash equilibrium is also a Nash equilibrium. An important and well-known result due to Kuhn ([16]), establishes that every extensive game (with a finite horizon) of perfect information has a subgame perfect Nash equilibrium in pure strategies. Closely related is the method of *backward induction*, which is essentially an algorithm providing subgame perfect Nash equilibria.

Each strategy profile corresponds with a collection of paths through the tree, each of which starts at a different internal node. In particular, a strategy profile determines a path connecting the root with a leaf. Similarly, strategies can be construed as subgraphs of (V, R). Another interesting subgraph results if one takes the union of a strategy profile s and the set of edges with the vertices possessed by a (sub-)set of players M as source. Intuitively, the significance of any such graph is that it reflects which outcomes a set of players can force to come about if they operate in coalition and the strategies of the other players are given. For the game of Example 2.1, this graph for Player 1 and the strategy profile s_{RLll} curbing Player 2's freedom of action is depicted in Figure 4.

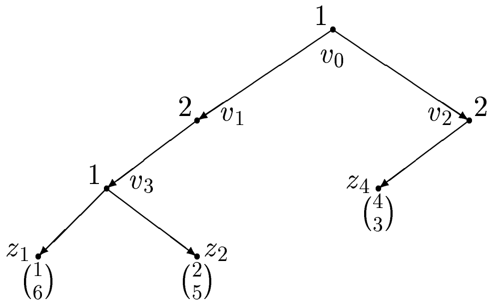

Figure 4.

To capture this notion formally we define for each strategy profile s and subset of players M a correspondence s_M on the vertices such that for all vertices $v \in V$:

$$s_M(v) \quad =_{df.} \quad \begin{cases} \{w : \ (v,w) \in R\} & \text{if } P(v) \in M, \\ \{s(v)\} & \text{otherwise}. \end{cases}$$

The correspondence s_M is obviously monotone in M, *i.e.*, $M' \subseteq M''$ implies $s_{M'} \subseteq s_{M''}$.

Each relation s_M induces a correspondence (set-valued function) on the vertices of the game, which value is a subset of the leaves of the tree. We define for each $s \in S$ and $M \subseteq N$, this correspondence $\hat{s}_M$ such that for each $v \in V$:

$$\hat{s}_M(v) \quad =_{df.} \quad \begin{cases} \{v\} & \text{if } v \text{ is a leaf,} \\ \bigcup \{\hat{s}_M(w) : \ w \in s_M(v)\} & \text{otherwise.} \end{cases}$$

We will write $\hat{s}_i$ for $\hat{s}_{\{i\}}$. The value the correspondence $\hat{s}_M$ takes at the root v_0 is the set of outcomes the players in M can force to come about by cooperating in the game if the other players adhere to the strategy profile s. In our example, $s_1(v_0) = \{z_1, z_2, z_4\}$, where s represents s_{RLll}. Obviously, the more players in M the larger this set of forceable outcomes, *i.e.*, the monotonicity of s_M propagates to $\hat{s}_M$:

$$M \subseteq M' \quad \text{implies} \quad \hat{s}_M \subseteq \hat{s}_{M'}.$$

The following fact relates notations and will prove to be particularly convenient.

Fact 2.1. Let s be a strategy profile of some extensive game G with v and v' vertices therein and M a subset of its players. Then:

$$v \in \hat{s}_M(v') \quad \text{iff} \quad \text{for some } s' \in S\text{:} \quad \hat{s}^M_{s'}(v') = v.$$

Obviously in particular $\hat{s}_\varnothing(v) = \{\hat{s}(v)\}$. The set of outcome nodes that can be reached by a strategy profile s with no player possibly deviating clearly contains as only the element the vertex that s determines as the unique outcome.

3. Describing and Reasoning about Extensive Games

Extensive games are based on trees. A similar remark applies to the players' preferences as they were defined over the vertices of the game tree in the previous section. Exploiting this relational structure, we propose a kind of multi-modal language to describe extensive games and reason about them. In particular, we will argue that such a language can express whether a strategy profile of an extensive game is a (subgame perfect) Nash equilibrium.

3.1. Syntax and Semantics

Our formal researches are conducted within propositional multi-modal logic. A propositional multi-modal language $\mathcal{L}(A, B)$ contains a non-empty but countable set of propositional variables A along with a countable set of labels B for monadic modalities. The formulas of $\mathcal{L}(A, B)$ are thus given by the following BNF-grammar, with $a \in A$ and $\beta \in B$:

$$\varphi \quad ::= \quad a \;\big|\; \neg\varphi \;\big|\; \varphi_0 \wedge \varphi_1 \;\big|\; [\beta]\varphi.$$

We assume the set of labels B to be the union of two disjoint sets and their Cartesian product, *i.e.*, $B = B_0 \cup B_1 \cup (B_0 \times B_1)$ with $B_0 \cap B_1 = \emptyset$. Moreover, B_0 will be assumed to be non-empty and finite. Multi-modal languages $\mathcal{L}(A, B)$ with B structured thus we will refer to as *multi-modal matrix languages*.

Extensive games are taken as the basis of the frames any such multi-modal language describes. Truth-value assignments to the propositional variables at each vertex takes care of the interpretation of the propositional variables and the Boolean connectives are given their conventional interpretation. The labels in B_0 go proxy for the players of a game. For each $\beta \in B_0$, the accessibility relation R_β runs along the preference relation of one of the players of the game. This gives rise to the intuitive reading of $[\beta]\varphi$ as "*φ holds in all states at least as preferable to i as the present one,*" where i is the player associated with the label β. For convenience, the labels in B_0 are also called *player labels*. In contrast, the labels in B_1 stand for strategy profiles of the game $\mathcal{L}(A, B)$ aims to describe. For each label $\beta \in B_1$ the accessibility relation R_β is defined as (the graph of) the function $\hat{s}$, where s is the strategy profile associated with β. As such R_β relates vertices to leaves only, being reflexive at the latter. Intuitively, $[\beta]\varphi$ then reads " *if, starting in the state of evaluation, all players choose their strategies as prescribed in s, the game ends in a situation in which φ holds,*" where s is the strategy profile associated with β. Finally, let i be the player associated with the label β in B_0 and s the strategy profile associated with the label β' in B_1. Then, the accessibility relation $R_{(\beta,\beta')}$ connects vertices v to the leaf nodes in $\hat{s}_i\,(v)$. With $\hat{s}_i\,(v)$ collecting all the terminal nodes that player i can force to come about provided that the other players adhere to the strategy profile s, $[(\beta, \beta')]\varphi$ obtains the informal interpretation of "*φ holds in all outcome states that can be reached if at most player i deviates from s.*"

The frames and models for the multi-modal languages are also of a special kind. Rather than taking into account all relational structures, the formal semantics is defined on frames that are structurally closely related to extensive games. A multi-modal language $\mathcal{L}(A, B)$ is said to be *feasible* for an extensive game G if B_0 is finite and for each of the labels in B_0 there is a player in the game and, similarly, for each label in B_1 there is a strategy profile s in G. Usually we will assume that B_0 is a subset of the players of the game, and B_1 a subset of strategy profiles.[1]

[1]The reason for not taking B_1 to be the whole set of strategy profiles is mainly of a technical nature. Requiring a lan-

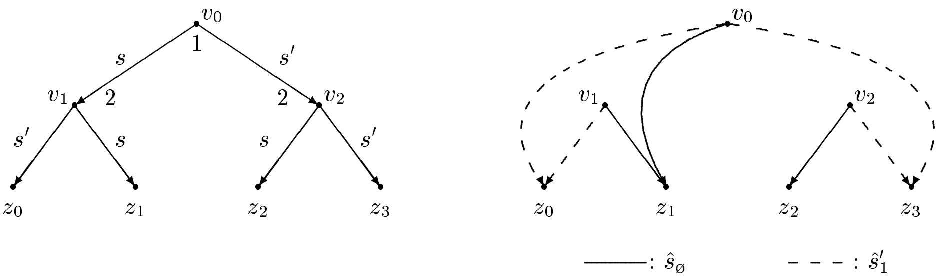

Figure 5. Transformation of an extensive game (left) to a game-frame (right) with respect to two strategy profiles s and s' and their corresponding accessibility relations $R_{\hat{s}_\varnothing}$ and $R_{\hat{s}'_1}$. In the righthand figure the reflexive arrows at the leaves are omitted.

The games for which $\mathcal{L}(A, B)$ is feasible constitute the class of so-called $\mathcal{L}(A, B)$-*feasible games* and is denoted by $\mathcal{G}_B$. For $\mathcal{L}(A, B)$ a multi-modal language, we associate with each extensive game G in $\mathcal{G}_B$ a frame as laid down formally in the following definition. The notion of a *game-model* for $\mathcal{L}(A, B)$ is then introduced much in the usual fashion.

Definition 3.1. (Game-frames and game-models)
Let G be an $\mathcal{L}(A, B)$-feasible extensive game given by $\big(V, R, N, P, \{\rho_i\}_{i\in N}\big)$. Define the frame $\mathfrak{F}_G$ for $\mathcal{L}(A, B)$ as the tuple $\big(V, \{R_\beta\}_{\beta\in B}\big)$, where for each $i \in B_0$ and each $s \in B_1$, the accessibility relations R_i, R_s and $R_{(i,s)}$ are such that for all v, v' in V:

$$\begin{aligned} vR_iv' &\quad \text{iff} \quad \big(v, v'\big) \in \rho_i \\ vR_sv' &\quad \text{iff} \quad v' \in \hat{s}_\varnothing\,(v) \\ vR_{(i,s)}v' &\quad \text{iff} \quad v' \in \hat{s}_i\,(v)\,. \end{aligned}$$

On the basis of these definitions we will generally write $[\hat{s}_\varnothing]\varphi$ and $[\hat{s}_i]\varphi$ for $[s]\varphi$ and $[(i, s)]\varphi$, respectively. The class of *game-frames* for $\mathcal{L}(A, B)$ is then given by $\{\mathfrak{F}_G : \ G \in \mathcal{G}_B\}$. A *game-model* $\mathfrak{M}$ for $\mathcal{L}(A, B)$ is a pair $(\mathfrak{F}, I)$ where $\mathfrak{F} = \big(V, \{R_\beta\}_{\beta\in B}\big)$ is a game-frame for $\mathcal{L}(A, B)$ and I a function assigning to each vertex in V a subset of propositional variables in A, *i.e.*, $I\colon V \to 2^A$. Figure 5 illustrates the construction of a game-frame from an extensive game.

This culminates in a standard interpretation of formulas in models.

$$\begin{aligned} &\mathfrak{M}, v \Vdash a && \text{iff} \quad a \in I(v) \\ &\mathfrak{M}, v \Vdash \neg\varphi && \text{iff} \quad \mathfrak{M}, v \nVdash \varphi \\ &\mathfrak{M}, v \Vdash \varphi \wedge \psi && \text{iff} \quad \mathfrak{M}, v \Vdash \varphi \ \text{ and } \ \mathfrak{M}, v \Vdash \psi \\ &\mathfrak{M}, v \Vdash [\beta]\varphi && \text{iff} \quad \text{for all } v' \in V \text{ such that } vR_\beta v' : \ \mathfrak{M}, v' \Vdash \varphi. \end{aligned}$$

Then $\mathfrak{M} \Vdash \varphi$ denotes that for all vertices v in $\mathfrak{M}$, it is the case that $\mathfrak{M}, v \Vdash \varphi$. We will use $\Gamma \vDash_{\mathsf{F}} \varphi$ to signify that in all vertices v of all models on a *game-frame* for $\mathcal{L}(A, B)$, if $\mathfrak{M}, v \Vdash \gamma$ for all γ in Γ then

guage $\mathcal{L}(A, B)$ to contain a label in B_1 for each strategy profile in the game to be described would jeopardize the construction of the model $\mathfrak{M}_\Gamma$ in the completeness proof, below. Yet, with respect to strategy profiles that are represented by a label in B_1 the language does not lose expressive power with respect to whether it is a (subgame perfect) Nash equilibrium.

also $\mathfrak{M}, v \Vdash \varphi$. Moreover, $\mathfrak{M}, v \Vdash \Gamma$ and $\mathfrak{M} \Vdash \Gamma$ abbreviate for all γ in Γ, $\mathfrak{M}, v \Vdash \gamma$ and for all γ in Γ, $\mathfrak{M} \Vdash \gamma$, respectively.

The modal semantics of multi-modal matrix languages is thus confined to models on game-frames, rather than that it pertains to *all* relational structures as is common practice in modal logic.

3.2. Characterizing Subgame Perfect Nash Equilibria

A strategy profile may or may not be a (subgame perfect) Nash equilibrium. If a strategy profile is a (subgame perfect) Nash equilibrium in an extensive game G, this fact is reflected in certain structural properties of the frame $\mathfrak{F}_G$. The aim of this section is to characterize these structural properties of a frame by means of a formula schema in a suitable multi-modal language. For G an $\mathcal{L}(A, B)$-feasible extensive game and s a strategy profile of G, what we are after is a formula schema $\vartheta(s)$ such that:

$$\mathfrak{F}_G \vDash \vartheta(s) \quad \text{iff} \quad s \text{ is a subgame perfect Nash equilibrium in } G.$$

It turns out that such a formula schema can be obtained as a special case of a schema well-known in standard modal correspondence theory. We say that a frame $\left(V, \{R_\beta\}_{\beta \in B}\right)$ for a modal language $\mathcal{L}(A, B)$ — not *per se* a multi-modal matrix language as introduced in the previous subsection — containing k, l, m and n as labels in B for $\mathcal{L}(A, B)$ has the (k, l, m, n)-***confluence property*** if:

$$\text{for all } v, w, x \in V: \; vR_k w \text{ and } vR_m x \; \text{ imply } \; \text{for some } y \in V: \; wR_l y \text{ and } xR_n y.$$

Here the labels k, l, m and n are *not* assumed to be necessarily distinct. The following fact, for a proof of which the reader be referred to [27], then holds:

Fact 3.1. (Confluence)
Let $\mathcal{L}(A, B)$ be a multi-modal language containing k, l, m and n as labels. Then the formula schema $\langle k\rangle[l]\varphi \to [m]\langle n\rangle\varphi$ characterizes frames for $\mathcal{L}(A, B)$ satisfying the (k, l, m, n)-confluence property.

If R_n is taken to be the identity relation on the set of vertices (k, l, m, n)-confluence reduces to the following property, which for obvious reasons we could dub (k, l, m)-***Euclidicity***:

$$\text{for all } v, w, x \in V: \; vR_k w \text{ and } vR_m x \; \text{ imply } \; wR_l x.$$

As a special case of Fact 3.1 we now obtain as a corollary the following fact, of which also the direct proof is elementary:

Corollary 3.1. For $\mathcal{L}(A, B)$ a multi-modal language containing k, l and m as labels, the formula schema $\langle k\rangle[l]\varphi \to [m]\varphi$ characterizes frames for $\mathcal{L}(A, B)$ satisfying (k, l, m)-Euclidicity.

By appropriately choosing k, l and m from the labels of a multi-modal language $\mathcal{L}(A, B)$ a strategy profile s being a subgame perfect best response for a player i in a game-frame $\mathfrak{F}$ can be characterized by the formula schema $\langle k\rangle[l]\varphi \to [m]\varphi$. Taking $\hat{s}_i$ for k, i for l and $\hat{s}_\varnothing$ for m, respectively, gives the desired result. Considering that this schema characterizes frames satisfying $(\hat{s}_i, i, \hat{s}_\varnothing)$-Euclidicity this makes informally sufficient sense:

$$\text{for all } v, v', v'' \in V: \; vR_{\hat{s}_i} v' \text{ and } vR_{\hat{s}_\varnothing} v'' \; \text{ imply } \; v'R_i v''.$$

In words this formula says that, if play commences at a vertex v, player i values the vertex v'' that the strategy profile s determines as an outcome at least as highly as any vertex v' that i can force to come about by unilaterally deviating from s. If this is the case, by deviating from s the player i will not be better of than by sticking to the strategy prescribed by s. The following proposition establishes this observation as an appropriate basis for the characterization of the game-theoretical property of a strategy profile being a subgame perfect response for a player in a game.

Proposition 3.1. Let $\mathcal{L}(A, B)$ be a multi-modal matrix language, let G be an extensive game in $\mathcal{G}_B$. For $i \in B_0$ and $s \in B_1$, then:

$$s \text{ is a subgame perfect best response for } i \text{ in } G \text{ iff } \mathfrak{F}_G \text{ is } (\hat{s}_i, i, \hat{s}_\emptyset)\text{-Euclidean.}$$

Proof:
For the left-to-right direction, assume the contrapositive, *i.e.*, that $\mathfrak{F}_G$ is *not* $(\hat{s}_i, i, \hat{s}_\emptyset)$-Euclidean. Then, there are vertices v, v' and v'' such that:

$$(a) \quad vR_{\hat{s}_i}v' \qquad (b) \quad vR_{\hat{s}_\emptyset}v'' \qquad (c) \quad \textit{not: } v'R_iv''.$$

By definition of $\mathfrak{F}_G$ these claims correspond to:

$$(a') \quad v' \in \hat{s}_i(v) \qquad (b') \quad v'' \in \hat{s}_\emptyset(v) \qquad (c') \quad (v', v'') \notin \rho_i.$$

With (a') and Fact 2.1 there is some s' such that $\hat{s}^i_{s'}(v) = v'$. Moreover, since $\hat{s}_\emptyset(v) = \{\hat{s}(v)\}$, also $\hat{s}(v) = v''$. Hence, with (c'), $(\hat{s}^i_{s'}(v), \hat{s}(v)) \notin \rho_i$, *i.e.*, s is *not* a subgame perfect Nash equilibrium in G.

For the right-to-left direction, assume that s is not a subgame perfect Nash equilibrium in G. Then for some vertex v, some player i and some strategy profile s', $(\hat{s}^i_{s'}(v), \hat{s}(v)) \notin \rho_i$. By definition of $\mathfrak{F}_G$, however, both $vR_{\hat{s}_i}\hat{s}^i_{s'}(v)$ and $vR_{\hat{s}_\emptyset}\hat{s}(v)$. It follows that $\mathfrak{F}_G$ is not $(\hat{s}_i, i, \hat{s}_\emptyset)$-Euclidean. □

Putting things together we obtain the following theorem, which lays down the results we were after. Recall that a player is called interested if she values some vertices strictly higher than other vertices.

Theorem 3.1. Let $\mathcal{L}(A, B)$ be a multi-modal matrix language and let $G \in \mathcal{G}_B$. Assume that M be a subset of B_0 containing a label for each *interested* player in G. Then, for i a player and s a strategy profile in G both represented by a label in B:[2]

(*i*)	s is a best response for i in G	iff	$\mathfrak{F}_G, v_0 \Vdash \langle\hat{s}_i\rangle[i]\varphi \to [\hat{s}_\emptyset]\varphi$
(*ii*)	s is a s.p. best response for i in G	iff	$\mathfrak{F}_G \Vdash \langle\hat{s}_i\rangle[i]\varphi \to [\hat{s}_\emptyset]\varphi$
(*iii*)	s is a Nash equilibrium in G	iff	$\mathfrak{F}_G, v_0 \Vdash \bigwedge_{i \in M} (\langle\hat{s}_i\rangle[i]\varphi \to [\hat{s}_\emptyset]\varphi)$
(*iv*)	s is a s.p. Nash equilibrium in G	iff	$\mathfrak{F}_G \Vdash \bigwedge_{i \in M} (\langle\hat{s}_i\rangle[i]\varphi \to [\hat{s}_\emptyset]\varphi)$.

[2]Here $\langle\hat{s}_i\rangle[i]\varphi \to [\hat{s}_\emptyset]\varphi$ and $\bigwedge_{i \in B_0} (\langle\hat{s}_i\rangle[i]\varphi \to [\hat{s}_\emptyset]\varphi)$ denote *formula schemas*, rather than formulas. Furthermore, 's.p.' abbreviates 'subgame perfect'.

Proof:
As all claims rest on much the same principles, we only present the proof for (*iv*) here. For the left-to-right direction, first assume that the formula schema $\bigwedge_{i\in M}\big(\langle\hat{s}_i\rangle[i]\varphi \rightarrow [\hat{s}_\emptyset]\varphi\big)$ is not valid in $\mathfrak{F}_G$. Hence, for some player i and some formula φ we have $\mathfrak{F}_G \not\models \langle\hat{s}_i\rangle[i]\varphi \rightarrow [\hat{s}_\emptyset]\varphi$. Consequently, the formula schema $\langle\hat{s}_i\rangle[i]\varphi \rightarrow [\hat{s}_\emptyset]\varphi$ is not valid in $\mathfrak{F}_G$ either. In virtue of Corollary 3.1, then, $\mathfrak{F}_G$ does not satisfy $(\hat{s}_i, i, \hat{s}_\emptyset)$-Euclidicity. With Proposition 3.1, then s is not a subgame perfect best response for i, and, *a fortiori*, neither a subgame perfect Nash equilibrium.

For the opposite direction, assume for some vertex v, some player i and some strategy profile s', that $\big(\hat{s}^i_{s'}(v), \hat{s}(v)\big) \notin \rho_i$. Observe that this renders i an interested player. An easy little inductive argument, which we will leave to the reader, establishes that $\hat{s}^i_{s'}(v) \in \hat{s}_i(v)$. Hence, by definition of $\mathfrak{F}_G$ both $vR_{\hat{s}_i}\hat{s}^i_{s'}(v)$ and $vR_{\hat{s}_\emptyset}\hat{s}(v)$. It follows that $\mathfrak{F}_G$ is not $(\hat{s}_i, i, \hat{s}_\emptyset)$-Euclidean. Hence, the formula schema $\langle\hat{s}_i\rangle[i]\varphi \rightarrow [\hat{s}_\emptyset]\varphi$ is not valid on $\mathfrak{F}_G$ and *a fortiori* neither the formula schema $\bigwedge_{i\in M}\big(\langle\hat{s}_i\rangle[i]\varphi \rightarrow [\hat{s}_\emptyset]\varphi\big)$. This concludes the proof. □

For each label $i \in B_0$ and each label $s \in B_1$ we refer to the axiom schema $\langle\hat{s}_i\rangle[i]\varphi \rightarrow [\hat{s}_\emptyset]\varphi$ by $\mathit{5}_{s,i}$ and to the axiom schema $\bigwedge_{i\in M}\big(\langle\hat{s}_i\rangle[i]\varphi \rightarrow [\hat{s}_\emptyset]\varphi\big)$ by $\mathit{5}^M_s$.

4. Axiomatization

This section concerns a number of axiom schemas for multi-modal matrix languages. In terms of these we introduce particular extensive game logics, F, $\mathsf{F5}_{s,i}$ and $\mathsf{F5}^M_s$, and discuss a number of completeness results, in particular that of F with respect to the class of *all* game-models. For the technical details of the completeness proofs the reader is referred to their exposition in the appendix to this paper.

4.1. The Axioms

For a multi-modal matrix language $\mathcal{L}(A, B)$ we have the following axioms. We assume β, β', β'', and β''' to range over the whole of B, β_0 and β'_0 over B_0, and β_1 and β'_1 over B_1.

Taut.: any classical tautology.

$$K: \quad [\beta]\big(\varphi \rightarrow \psi\big) \rightarrow \big([\beta]\varphi \rightarrow [\beta]\psi\big)$$

$$T_{\beta_0}: \quad [\beta_0]\varphi \rightarrow \varphi$$

$$4_{\beta_0}: \quad [\beta_0]\varphi \rightarrow [\beta_0][\beta_0]\varphi$$

$$D!_{\beta_1}: \quad [\beta_1]\varphi \leftrightarrow \langle\beta_1\rangle\varphi$$

$$F1_{(\beta_0,\beta_1),\beta_1}: \quad [(\beta_0, \beta_1)]\varphi \rightarrow [\beta_1]\varphi$$

$$F2_{(\beta_0,\beta_1),(\beta'_0,\beta'_1)}: \quad [(\beta_0, \beta_1)]\big([(\beta'_0, \beta'_1)]\varphi \leftrightarrow \varphi\big)$$

$$F3_{\beta,\beta',\beta'',\beta''',\beta_0}: \quad [\beta][\beta']\big([\beta_0]\varphi \rightarrow \psi\big) \vee [\beta''][\beta''']\big([\beta_0]\psi \rightarrow \varphi\big).$$

The logic F is closed under the rules of *modus ponens* (*MP*) and *necessitation* (*Nec.*):

$$MP: \quad \frac{\varphi \rightarrow \psi \qquad \varphi}{\psi} \qquad\qquad Nec.: \quad \frac{\varphi}{[\beta]\varphi}.$$

The Hilbert-style axiom system given by these axioms and rules we will refer to as the normal modal logic F. Any multi-modal logic containing F we will refer to as an *extensive game logic*. Accordingly, F is the smallest extensive game logic, in a similar manner as K is the smallest normal modal logic.

Definition 4.1. (Extensive game logics)
An *extensive game logic* F for a multi-modal matrix language $\mathcal{L}(A,B)$ is any set of formulas of $\mathcal{L}(A,B)$ closed under *MP* and *Nec.* and containing all instances of the axiom schemes of *Taut.*, *K*, $\mathit{4}_{\beta_0}$, T_{β_0}, $D!_{\beta_1}$, $F1_{(\beta_0,\beta_1),\beta_1}$, $F2_{(\beta_0,\beta_1),(\beta_0',\beta_1')}$ and $F3_{\beta,\beta',\beta'',\beta''',\beta_0}$. We will write $\Gamma \vdash_\Lambda \varphi$ if there exists a derivation of φ from the theory Γ in an extensive game logic Λ, as usual. The smallest extensive game logic we will refer to by F.

At the conclusion of this section we will come to review also the stronger extensive game logics than F, *viz.*, the logics $\mathsf{F5}_{s,i}$ and $\mathsf{F5}_s^M$. The former has $5_{s,i}$, for fixed s and i, as an additional axiom and the latter 5_s^M for a particular subset M of β_0.

Within the setting of extensive games and the intended interpretation of the multi-modal matrix languages, the axioms *K* through $F3_{\beta,\beta',\beta'',\beta''',\beta_0}$ have quite intuitive readings. The axioms *Taut.* and *K* along with the two rules for *modus ponens* (*MP*) and necessitation (*Nec.*) guarantee extensive game logics to be normal logics. With the accessibility relations for the modal operators with labels in B_0 running over the preferences of players in an extensive game, T_{β_0} and $\mathit{4}_{\beta_0}$ warrant the players' preferences to be reflexive and transitive. The axiom $F3_{\beta,\beta',\beta'',\beta''',\beta_0}$ is sound in virtue of the players' preference relations being connected. Within the more comprehensive setting of general Kripke frames connectivity of a relation is not characterizable by a formula scheme. That $F3_{\beta,\beta',\beta'',\beta''',\beta_0}$ nevertheless succeeds in doing so here is due to the fact that the semantics is restricted to the class of game-frames, for which we may assume some additional structure. Observe that in virtue of *Taut.*, T_{β_0} and $F3_{\beta,\beta',\beta'',\beta''',\beta_0}$, we can derive the following axiom in each extensive game logic:

$$F4_{\beta,\beta',\beta_0}: \quad [\beta]\big([\beta_0]\varphi \to \psi\big) \vee [\beta']\big([\beta_0]\psi \to \varphi\big).$$

The labels in B_1 represent strategy profiles and in particular their outcome functions. The accessibility relation R_{β_1} connects, for some strategy profile s, any vertices v and v' such that $v' = \hat{s}\,(v)$ and as such is the graph of a function. Hence, Axiom $D!_{\beta_1}$. Axiom $F1_{(\beta_0,\beta_1),\beta_1}$ characterizes the inclusion of the accessibility relation R_{β_1} in $R_{(\beta_0,\beta_1)}$. The intuition behind this lies in the observation that any outcome that is determined by a strategy profile can also be reached if one of the players has the option to deviate from that strategy profile; the player in question may choose to adhere to the strategy profile after all. The labels in $B_0 \times B_1$ represent the correspondences $\hat{s}_i$ for strategy profiles s and players i. The value of any such correspondence is a set of leaf nodes, from each of which only the leaf itself can be reached. It is exactly this fact that $F2_{(\beta_0,\beta_1),(\beta_0',\beta_1')}$ reflects. Observe that as a consequence of $F2_{(\beta_0,\beta_1),(\beta_0',\beta_1')}$, $D!_{\beta_0}$ and $F1_{(\beta_0,\beta_1),\beta_1}$ we can derive the following more general axiom schema, in which both β and β' range over labels in either B_1 or in $B_0 \times B_1$:

$$F5_{\beta,\beta'}: \quad [\beta]\,\big([\beta']\varphi \leftrightarrow \varphi\big)\,.$$

Note that the scheme $F5_{\beta,\beta'}$ does not hold in general if β or β' are in B_0.

The cogency of these informal remarks are vindicated in the following proposition, which formally establishes the soundness on game-frames of the axioms in question.

Proposition 4.1. (Soundness)
The axioms *K* though $F3_{\beta,\beta',\beta_0}$ in a multi-modal language $\mathcal{L}(A,B)$ are valid on all of $\mathcal{L}(A,B)$'s game-frames.

Proof:
For ordinary multi-modal frames the axioms T_β and 4_β characterize reflexivity and transitivity of R_β, respectively. Similarly $D!_\beta$ characterizes functionality of R_β and $F1_{\beta,\beta'}$ the inclusion of $R_{\beta'}$ in R_β. The axiom schema $F2_{\beta,\beta'}$ characterizes frames in which $R_{\beta'}$ is the identity at every vertex v that is reachable by R_β, *i.e.*, frames for which:

$$\text{for all } v, v' \in V: \quad vR_\beta v' \ \text{ implies } \ \text{for all } v'' \in V: v'R_{\beta'}v'' \ \text{ iff } \ v' = v''.$$

Finally, $F3_{\beta,\beta',\beta'',\beta''',\beta_0}$ — $[\beta][\beta']\big([\beta_0]\varphi \to \psi\big) \vee [\beta''][\beta''']\big([\beta_0]\psi \to \varphi\big)$ — characterizes frames in which any two vertices v and v' are comparable with respect to R_{β_0} if the one is reachable from from some third vertex v'' *via* $R_\beta \circ R_{\beta'}$ and the other from the same vertex v'' *via* $R_{\beta''} \circ R_{\beta'''}$. *I.e.*, more formally, $F3_{\beta,\beta',\beta''}$ characterizes frames in which for all vertices v, v', v'':

$$\text{if} \quad \text{for some } w', w'': vR_\beta w' R_{\beta'} v' \ \text{ and } \ vR_{\beta''} w'' R_{\beta'''} v'' \quad \text{then} \quad v' R_{\beta_0} v'' \ \text{ or } \ v'' R_{\beta_0} v'.$$[3]

All of the above are results of elementary modal correspondence theory.

Since the game-frames for $\mathcal{L}(A,B)$ are special cases of ordinary multi-modal frames and the connectives obtain their usual Boolean interpretations *Taut.*, *K*, *MP* and *Nec.* hold without ado and it suffices to show that the properties T_{β_0} through $F3_{\beta,\beta',\beta''}$ characterize in ordinary frames are satisfied in the game-frames for $\mathcal{L}(A,B)$.

The players' preferences were assumed to be reflexive, transitive and connected and *a fortiori* so are R_{β_0} for each $\beta_0 \in N$. This takes care of the soundness of T_{β_0}, 4_{β_0} and $F3_{\beta,\beta',\beta'',\beta''',\beta_0}$. Strategy profiles determine a unique leaf node as outcome. Formally, for each strategy profile s and each vertex v, $\hat{s}_\emptyset(v) = \{\hat{s}(v)\}$, which renders functional the accessibility relation R_{β_1} for each β_1. Hence, $D!_{\beta_1}$ is valid in game-frames as well. In virtue of the monotonicity of $\hat{s}_M$ (*cf.*, page 289), we have in particular that $\hat{s}_\emptyset \subseteq \hat{s}_i$. Hence, also $R_{\beta_1} \subseteq R_{(\beta_0,\beta_1)}$ for all $\beta_0 \in N$ and $\beta_1 \in S$. The validity of $F1_{(\beta_0,\beta_1),\beta_1}$ follows. Finally, for $F2_{(\beta_0,\beta_1),(\beta_0,\beta_1)}$ it suffices to show that for all strategy profiles s and s' and for all vertices v and v' in a game-tree, $v' \in \hat{s}_i(v)$ implies $\hat{s}'_j(v') = \{v'\}$. Merely observe that in general $\hat{s}_i(v) \subseteq Z$ and that for all leaves $z \in Z$ we have that $\hat{s}'_j(z) = \{z\}$ by definition. □

For easy reference we have collected all axioms that have so far been dealt with in Table 1, where the labels are chosen in such a way as to reflect their intended game-theoretical readings as suggested in Definition 3.1, above. *I.e.*, typical elements of B_0, B_1 and $B_0 \times B_1$ are taken to be, respectively, i, $\hat{s}_\emptyset$ and $\hat{s}_i$.

4.2. Completeness

This section concerns some completeness results for the extensive game logics F, $\mathsf{F5}_{s,i}$ and $\mathsf{F5}_s^N$ in a multi-modal matrix language $\mathcal{L}(A,B)$. The extensive game logic F is complete with respect to the class

[3]This property is a close multi-modal relative of that of *piecewise connectedness*. A Kripke frame is said to be piecewise connected if for all vertices v, v' and v'', vRv' and vRv'' implies either $v'Rv''$ or $v''Rv'$.

Taut.:	any classical tautology.
K:	$[\beta]\big(\varphi \to \psi\big) \to \big([\beta]\varphi \to [\beta]\psi\big)$
T_i:	$[i]\varphi \to \varphi$
4_i:	$[i]\varphi \to [i][i]\varphi$
$5_{s,i}$:	$\langle \hat{s}_i \rangle [i]\varphi \to [\hat{s}_\varnothing]\varphi$
5_s^M:	$\bigwedge_{i\in M} \big(\langle \hat{s}_i \rangle [i]\varphi \to [\hat{s}_\varnothing]\varphi\big)$
$D!_{\hat{s}_\varnothing}$:	$[\hat{s}_\varnothing]\varphi \leftrightarrow \langle \hat{s}_\varnothing \rangle \varphi$
$F1_{\hat{s}_i,\hat{s}_\varnothing}$:	$[\hat{s}_i]\varphi \to [\hat{s}_\varnothing]\varphi$
$F2_{\hat{s}_i,\hat{s}'_j}$:	$[\hat{s}_i]\big([\hat{s}'_j]\varphi \leftrightarrow \varphi\big)$
$F3_{\beta,\beta',\beta'',\beta''',i}$:	$[\beta][\beta']\big([i]\varphi \to \psi\big) \vee [\beta''][\beta''']\big([i]\psi \to \varphi\big)$
$F4_{\beta,\beta',i}$:	$[\beta]\big([i]\varphi \to \psi\big) \vee [\beta']\big([i]\psi \to \varphi\big)$
$F5_{\hat{s}_X,\hat{s}'_Y}$:	$[\hat{s}_X]\,([\hat{s}'_Y]\varphi \leftrightarrow \varphi)$ where $X, Y \in N \cup \{\varnothing\}$

Table 1. List of axiom schemas for multi-modal matrix languages $\mathcal{L}(A, B)$, where $B = N \cup S \cup (N \times S)$ and β and its primed varieties range over B.

of all game-models, *i.e.*, for all formulas φ and theories Γ of a multi-modal matrix language $\mathcal{L}(A, B)$ we have:

$$\Gamma \vDash_{\mathsf{F}} \varphi \quad \text{implies} \quad \Gamma \vdash_{\mathsf{F}} \varphi.$$

Similar results can be obtained for the logics $\mathsf{F5}_{i,s}$ and $\mathsf{F5}_s^N$ (for fixed labels i and s, and N the whole set of player labels of the respective language). The former is complete with respect to the class of game-models built on games in which the strategy profile s contains a best response for player i. The latter is complete with respect to those models that are based on extensive games in which s is a subgame perfect Nash equilibrium, provided that N contains a label for each *interested* player, *i.e.*, a player valuing some possible outcomes strictly higher than others. In the appendix to this paper we give the full proofs of these claims. Still, some remarks as to their structure and consequences are in order.

In order to prove completeness of an extensive game logic Λ with respect to a certain class of game-models $\mathcal{C}$, it suffices, by a standard argument, to construct for each Λ-consistent theory Γ a game-model that satisfies Γ at some vertex and prove this model to be in the class $\mathcal{C}$. The main problem in proving completeness for an extensive game logic Λ is in the construction of a model $\mathfrak{M}_\Gamma^\Lambda$ for a theory Γ that is consistent in F. In particular, it should be guaranteed that this model $\mathfrak{M}_\Gamma^\Lambda$ is a *game-model* in the sense of Definition 3.1, *i.e.*, that the frame underlying $\mathfrak{M}_\Gamma^\Lambda$ is based on an extensive game. Although the axioms of F, $\mathsf{F5}_{i,s}$ and $\mathsf{F5}_s^N$ are all of a standard nature, it is not obvious, however, whether a canonical model that would result in a standard Henkin-style proof — *i.e.*, a model satisfying *each* Λ-consistent theory at the same time — is indeed based on a game-frame. A proof of this is likely to become complicated if not impossible because the structure of a canonical model is pretty much fixed.

A construction method, in which for each Λ-consistent theory Γ separately a model $\mathfrak{M}_\Gamma^\Lambda$ satisfying Γ is defined, affords more control over the structure of the model to be built. In particular, the process of constructing the model $\mathfrak{M}_\Gamma$ for a theory Γ can be made go hand in hand with the construction of an extensive game $\boldsymbol{G}_\Gamma$ underlying $\mathfrak{M}_\Gamma$. As such this method of proof is more natural for our purposes.

For an F-consistent theory, any model satisfying Γ will do, provided it be a *game-model*. In the appendix we show how from any F-consistent theory, a game-model $\mathfrak{M}_\Gamma^\Lambda$ can be constructed that satisfies Γ. This construction can be extrapolated as to apply to extensive game logics in general, thus providing a game-model $\mathfrak{M}_\Gamma^\Lambda$ for each Λ-consistent theory Γ. For the completeness of $\mathsf{F5}_{i,s}$ and $\mathsf{F5}_s^N$ it has additionally to be proved that the models this construction gives are in the appropriate class of models. *I.e.*, for $\mathsf{F5}_{i,s}$ it has to be shown that for each $\mathsf{F5}_{i,s}$-consistent theory Γ the model $\mathfrak{M}_\Gamma^{\mathsf{F5}_{i,s}}$ is based on an extensive game in which the strategy profile s contains a best response for player i. In the case of $\mathsf{F5}_s^N$, similarly, one should show that, in the extensive game underlying a model $\mathfrak{M}_\Gamma^{\mathsf{F5}_i^N}$, for each interested player for a $\mathsf{F5}_s^N$-consistent theory Γ, there be a label in N, and, moreover, the strategy profile denoted by the label s is a subgame perfect Nash equilibrium.

The appendix describes in detail the construction of the game-model $\mathfrak{M}_\Gamma^\Lambda$ for each Λ-consistent theory Γ. Here, we confine ourselves to pointing out that in order to guarantee that $\mathfrak{M}_\Gamma^\Lambda$ be a *game-model*, in the same process also an extensive game $\boldsymbol{G}_\Gamma$ is defined which underlies $\mathfrak{M}_\Gamma^\Lambda$, *i.e.*, the game-model $\mathfrak{M}_\Gamma^\Lambda$ is defined as a model on the game-frame $\mathfrak{F}_{\boldsymbol{G}_\Gamma}$. The model $\mathfrak{M}_\Gamma^\Lambda$ is then shown to satisfy the theory Γ at the root.

The completeness of F with respect to all game-models then follows by a standard argument. For, assume for an arbitrary theory Γ and an equally arbitrary formula φ that $\Gamma \nvdash_{\mathsf{F}} \varphi$. Then, $\Gamma \cup \{\neg\varphi\}$ is F-consistent. Consequently, the game-model $\mathfrak{M}_{\Gamma\cup\{\neg\varphi\}}^{\mathsf{F}}$, as defined as above, satisfies the theory $\Gamma \cup \{\neg\varphi\}$ at the root, thus establishing that $\Gamma \nvDash_{\mathsf{F}} \varphi$. With soundness of F already having been proved, we have the following theorem.

Theorem 4.1. (Completeness of F)
Let Γ be a theory and φ a formula in a multi-modal matrix language $\mathcal{L}(A, B)$. Then:

$$\Gamma \vdash_{\mathsf{F}} \varphi \quad \text{iff} \quad \Gamma \vDash_{\mathsf{F}} \varphi.$$

For the completeness of $\mathsf{F5}_{s,i}$, it should additionally be proved that for each $\mathsf{F5}_{s,i}$-consistent theory Γ in the extensive game $\boldsymbol{G}_\Gamma$, on which the game model $\mathfrak{M}_\Gamma^{\mathsf{F5}_{s,i}}$ is based, the strategy profile (denoted by) s contains a subgame perfect best response for player (denoted by) i. In the appendix we show that this is actually the case, giving rise to the following theorem.

Theorem 4.2. The logic $\mathsf{F5}_{i,s}$ is sound and complete with respect to the class of game-frames built on extensive games in which s is a subgame perfect best response for player i.

Similarly, if Γ is an $\mathsf{F5}_s^N$-consistent theory, for all interested players of the game $\boldsymbol{G}_\Gamma$ underlying $\mathfrak{M}_\Gamma^{\mathsf{F5}_s^i}$ there is a label in N. Moreover, the strategy profile represented by s is then a subgame perfect Nash-equilibrium in $\boldsymbol{G}_\Gamma$. Hence also:

Theorem 4.3. The logic $\mathsf{F5}_s^N$ in $\mathcal{L}(A, B)$ is sound and complete with respect to the class of game-frames built on games in which s is a subgame perfect Nash equilibrium and in which there is a label in N for each interested player.

As an immediate consequence of these completeness results and the fact that in any derivation of a formula φ from a theory Γ in an extensive game logic Λ only a finite number of formulas can occur, we have that each of the logics F, $\mathsf{F5}_{s,i}$ and $\mathsf{F5}_s^N$ is *compact*. *I.e.*, for each theory Γ and each formula φ, if $\Gamma \vDash_\Lambda \varphi$ then there is a *finite* subtheory Γ_e of Γ such that $\Gamma_e \vDash_\Lambda \varphi$.

For Γ a Λ-consistent theory in a language $\mathcal{L}(A, B)$, the game $\boldsymbol{G}_\Gamma$ has in general some noteworthy properties. In particular, the depth of the game $\boldsymbol{G}_\Gamma$ — *i.e.*, the length of the longest path in the game-tree connecting the root to a leaf — does not exceed the number of player labels in $\mathcal{L}(A, B)$ plus two, *i.e.*, $\|N\| + 2$. Moreover, the players are assumed to play in a fixed order, and on each path in the game-tree from the root to a leaf, each player represented by a label in N moves at most once and any other player at most twice. Also, the number of players in each game $\boldsymbol{G}_\Gamma$ is always one greater than the number of labels in N.

Corollary 4.1. Let G be an extensive game of perfect information G with N the set of players and let Γ be a theory in a feasible multi-modal matrix language $\mathcal{L}(A, B)$ with N the set of player labels. Assume further that Γ is satisfiable in a model on $\mathfrak{F}_G$. Then there is an extensive game of perfect information G' a game-model on which also satisfies Γ and which game tree has a maximal depth of $\|N\| + 2$ and which players number $\|N\| + 1$. Moreover, in each play of G' each player represented by a label in N moves at most once and any other player at most twice.

This corollary says, roughly, that one can confine one's attention to games of a limited depth when studying finite extensive games with respect to properties of extensive games expressible in the respective multi-modal matrix language.

5. Characterizing Nash Equilibria in PDL

In Section 3, we argued that the structural dependencies that obtain between the players' preferences and their strategies when a strategy profile is a (subgame perfect) Nash equilibrium can suitably be characterized by means of multi-modal matrix languages. Some of the labels of a multi-modal matrix language $\mathcal{L}(A, B)$ represent a strategy profile s and are interpreted as the graph of $\hat{s}_\emptyset$. These labels, however, have no further internal structure, and consequently in order to evaluate a formula of the form $[\hat{s}_\emptyset]\varphi$, one needs to calculate, independently of the semantics, the value of $\hat{s}_\emptyset(v)$ in the game under scrutiny. This is a plain game-theoretic task. A similar remark applies to the evaluation of formulas of the form $[\hat{s}_i]\varphi$ in a vertex v, which requires the calculation of the value of $\hat{s}_i(v)$. Much of the game-theoretical burden has thus been put on the *interpretation* of the multi-modal matrix languages on the models. Putting it slightly differently, the transformation of a game to a game-frame requires extensive game-theoretical reasoning on the meta-level.

For an illustration of this point consider once again the game of Example 2.1, above. Let $\mathcal{L}(A, B)$ be a feasible multi-modal matrix language containing a label s for the strategy profile s_{RLll}. In order to evaluate, *e.g.*, the formula of the form $[\hat{s}_2]\varphi$ at a state v in a model on the corresponding frame, one should investigate whether φ holds in all vertices in $\hat{s}_2(v)$. The relation $R_{\hat{s}_2}$, however, is taken to be semantically primitive and to establish that $v_0 R_{\hat{s}_2} z_5$ but that not $v_3 R_{\hat{s}_2} z_2$ the semantics is of no further help; these facts have to be obtained independently at the meta-level of reasoning.

In the dynamic language of *Propositional Dynamic Logic* (PDL) the set of labels have a richer structure giving rise to a highly expressive modal logic. Exploiting this structure and expressive power some of the game-theoretical burden can be shifted from the informal meta-level of reasoning about the model and the interpretation to the logic itself. We will find that the relations corresponding to $\hat{s}_i$ and $\hat{s}_\emptyset$ are the accessibility relations associated with labels denoting complex programs which allow for further semantical analysis. Also the way a frame for an appropriate dynamic language is manufactured from an extensive game will be more direct and will preserve more of the treelike structure of an extensive game than was the case for multi-modal matrix languages (for an illustration of this point compare Figure 5, above, and Figure 6, below).

This section concerns a class of two-sorted multi-modal languages $\mathcal{L}(A, B)$, where B is the union of two disjoint sets B_0 and Π_{B_1}, where Π_{B_1} denotes the set of PDL-programs over a set B_1 of atomic programs. Also this set of atomic programs B_1 we assume to be the union of two disjoint sets B_{10} and B_{11}. Consequently, the formulas φ and programs π of such a language $\mathcal{L}(A, B)$ — which we will call *dynamic multi-modal languages* — are given by the following BNF-grammar, with $a \in A$, $\beta_0 \in B_0$ and $\beta_1 \in B_1$:

$$\begin{array}{lcl} \varphi & ::= & a \mid \neg\varphi \mid \varphi_0 \wedge \varphi_1 \mid [\pi]\varphi \mid [\beta_0]\varphi \\ \pi & ::= & \beta_1 \mid \pi_0;\pi_1 \mid \pi_0 \cup \pi_1 \mid \pi^* \mid \varphi?. \end{array}$$

Extensive games are again used as the basis for the models on which such languages are interpreted. The propositional connectives and the program operators obtain their usual informal readings of negation ($\neg$), conjunction ($\wedge$), sequential composition (;), non-deterministic choice ($\cup$), iteration (*) and test (?). We also have the usual abbreviations, in particular that of "$\texttt{while}\ \varphi\ \texttt{do}\ \pi\ \texttt{od}$" for "$(\varphi?;\pi)^*;\neg\varphi$". The labels in B_0 go proxy for the players of a game, giving rise to the informal reading of $[i]\varphi$ as "*φ holds in all states at least as preferable to i as the state of evaluation*", as before. The atomic programs in B_1 are interpreted as a subset of the edges of the game-tree. Each atomic program $\beta \in B_{10}$ is associated with a player i and runs along those edges (v, v') of which v is assigned to the player i. So, intuitively, $[\beta_{10}]\varphi$, with β_{10} associated with player i, reads "*if i is to move, then φ holds at the next stage of the game no matter which strategy i decides to act upon.*" The atomic programs β_{11} in B_{11} are each associated with a strategy profile s. Informally, $[\beta_{11}]\varphi$ holds at a vertex v if φ holds at the next stage the game will be in if the strategy profile s is adhered to.

As before, we use the set of extensive games to define the class of frames and models for such a language. We will assume for each label in B_0 there to be a player in the game, and a one-one correspondence between B_0 and the set of labels B_{10}. Usually we will identify B_0 with a subset of players of the game $\mathcal{L}(A, B)$ means to describe. The atomic program thus associated with a player i will be denoted by $\pi(i)$. In a similar, fashion we will assume that for each label in B_{11} there is a different strategy profile of the game in question. For s a strategy profile, $\pi(s)$ denotes the label in B_{11} it is thus associated with. These conditions define *feasibility* of dynamic languages $\mathcal{L}(A, B)$ for extensive games. We now define formally:

Definition 5.1. (Dynamic game-frames and dynamic game-models)
Let $\mathcal{L}(A, B)$ a dynamic multi-modal language. Let G be an extensive game $\left(V, R, N, P, \{\rho_i\}_{i \in N}\right)$ for which $\mathcal{L}(A, B)$ is feasible. Define the *dynamic game-frame* $\mathfrak{F}_G^{\mathsf{PDL}}$ as the tuple $\left(V, \{R_\beta\}_{\beta \in B}\right)$, where for each $i \in B_0$, each $\pi(i) \in B_{10}$ and each $\pi(s) \in B_{11}$, the relations R_i, $R_{\pi(i)}$ and $R_{\pi(s)}$ are such that for

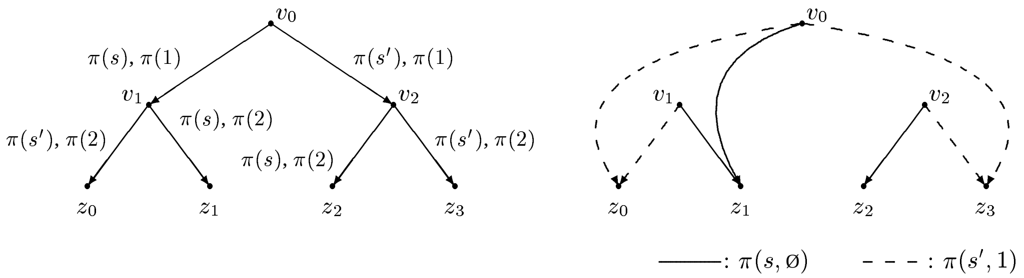

Figure 6. Transformation from the extensive game in Figure 5 to a dynamic game-frame (left) with respect to two strategy profiles s and s' and their corresponding atomic programs $\pi(s)$, $\pi(s')$, $\pi(1)$ and $\pi(2)$. The righthand figure shows the programs $\pi(s, \varnothing)$ and $\pi(s', 1)$. Note that these can be *"derived"* from the lefthand figure, whereas in the multi-modal framework these relations were primitive .

all $v, v' \in V$:

$$\begin{array}{rll} vR_i v' & \text{iff} & \big(v, v'\big) \in \rho_i \\ vR_{\pi(i)} v' & \text{iff} & P\left(v\right) = i \ \text{ and } \ \big(v, v'\big) \in R \\ vR_{\pi(s)} v' & \text{iff} & s\left(v\right) = v'. \end{array}$$

If formal rigor permits we will often omit the superscript PDL for esthetic reasons. A *dynamic game-model* $\mathfrak{M}$ for $\mathcal{L}(A, B)$ is defined as usual as a pair $(\mathfrak{F}, I)$, where $\mathfrak{F}$ is a dynamic game-frame for $\mathcal{L}(A, B)$ and I an interpretation function for the propositional variables in A. Figure 6 illustrates the transformation of an extensive game to a dynamic game frame.

The evaluation of formulas in a PDL-model is then as usual.

$$\begin{array}{lll} \mathfrak{M}, v \Vdash a & \text{iff} & a \in I(v) \\ \mathfrak{M}, v \Vdash \neg\varphi & \text{iff} & \mathfrak{M}, v \nVdash \varphi \\ \mathfrak{M}, v \Vdash \varphi \wedge \psi & \text{iff} & \mathfrak{M}, v \Vdash \varphi \ \text{ and } \ \mathfrak{M}, v \Vdash \psi \\ \mathfrak{M}, v \Vdash [\beta]\varphi & \text{iff} & \text{for all } v' \in V \text{ such that } vR_\beta v' : \ \mathfrak{M}, v' \Vdash \varphi. \end{array}$$

A PDL-model $\mathfrak{M}$ is said to be *regular* if program connectives ";", "$\cup$", "$*$" and "?" have their intuitive interpretations of *sequential composition*, *non-deterministic choice*, *iteration* and *test*, respectively, *i.e.*, if the following conditions are fulfilled:

$$\begin{array}{rcl} R_{\pi_1;\pi_2} & = & R_{\pi_1} \circ R_{\pi_2} \\ R_{\pi_1 \cup \pi_2} & = & R_{\pi_1} \cup R_{\pi_2} \\ R_{\pi^*} & = & (R_\pi)^* \\ R_{\varphi?} & = & \big\{(v, v) : \ \mathfrak{M}, v \vDash \varphi\big\}. \end{array}$$

Here $R_{\pi_1} \circ R_{\pi_2}$ denotes the relational composition of R_{π_1} and R_{π_2}, and $(R_\pi)^*$ is the transitive reflexive closure or ancestral of R_π. In the sequel we will assume PDL-models to be regular.

The important thing to observe in this definition is that the accessibility relations R_i, $R_{\pi(i)}$ and $R_{\pi(s)}$ can be read off from the extensive game specification almost immediately. In particular, the construction does not invoke the correspondences $\hat{s}_i$ and $\hat{s}_\varnothing$ for the interpretation of the atomic progams.

Theorem 3.1, above, showed that subgame perfect Nash equilibria are characterized in multi-modal matrix languages by the axiom schema $\bigwedge_{i\in N}\big(\langle\hat{s}_i\rangle[i]\varphi \rightarrow [\hat{s}_\varnothing]\varphi\big)$. The dynamic modal languages of this section do not possess the modalities $[\hat{s}_\varnothing]$ and $[\hat{s}_i]$ *explicitly*. However, for each dynamic modal language $\mathcal{L}(A,B)$ they can be defined implicitly as molecular PDL-programs. Let s be a label in B_{11} representing a strategy profile and let $\{i_0,\ldots,i_m\}$ be a subset of labels in B_0 denoted by M. Then, introduce the following abbreviation:

$$\pi(s,M) \quad =_{df.} \quad \texttt{while } \langle\pi(s)\rangle\top \texttt{ do } \pi(s)\cup\pi(i_0)\cup\ldots\cup\pi(i_m) \texttt{ od}.$$

We will write $\pi(s,i)$ for $\pi(s,\{i\})$. The idea is then that the program $\pi(s,i)$ performs the same task in PDL as the label $\hat{s}_i$ in the multi-modal languages, and, similarly, $\pi(s,\varnothing)$ is the dynamic counterpart of the multi-modal label $\hat{s}_\varnothing$. Construed as a program, $\pi(s,M)$ performs non-deterministically one of the programs $\pi(i)$ or $\pi(s)$, as long as $\pi(s)$ is enabled. Given the informal readings of the atomic programs $\pi(s)$ and $\pi(i)$ have in the dynamic game-models, $\pi(s,M)$ also allows for a rather more game-theoretical interpretation. The accessibility relation $R_{\pi(s,M)}$ connects vertices v with leaves z of the game tree, whenever z is a possible outcome state if play is commenced in v and the strategy profile s is adhered to by all players, with the possible exception of the players in M. Formally, the following proposition vindicates this intuitive interpretation.

Proposition 5.1. Let G denote the extensive game $\big(V,R,N,P,\{\rho_i\}_{i\in N}\big)$ and let be $\mathfrak{F}_G$ be a dynamic game-frame for a feasible dynamic multi-modal language $\mathcal{L}(A,B)$. Let furthermore $M\subseteq B_0$ and $s\in B_{11}$. Then for all vertices $v,v'\in V$:

$$vR_{\pi(s,M)}v' \quad \text{iff} \quad v'\in\hat{s}_M(v)\,.$$

Proof:
Consider an arbitrary model $\mathfrak{M}$ on $\mathfrak{F}_G$. Define the *height* of a vertex v in (V,R), denoted by *hgt* (v), inductively as:

$$\textit{hgt}\,(v) \quad =_{df.} \quad \begin{cases} 0 & \text{if } v \text{ is a leaf,} \\ 1+\textit{max}\big\{\textit{hgt}\,(v') : \ (v,v')\in R\big\} & \text{otherwise.}\end{cases}$$

The proof is then by induction on *hgt* (v).

For the basis assume *hgt* $(v)=0$. Then v is a leaf and we have $\hat{s}_M(v)=\{v\}$. Since v is a leaf there is no v' such that $s(v)=v'$ and accordingly, $\mathfrak{M},v\nVdash\langle\pi(s)\rangle\top$. Hence, the guard of $\pi(s,M)$ is not satisfied at v and $vR_{\pi(s,M)}v'$ if and only if $v'=v$, which proves the case.

For the induction step let *hgt* $(v)=n+1$. Then v is an internal node and by definition of a strategy profile there is some v' such that $s(v)=v'$, which makes that the guard of $\pi(s,M)$ is satisfied at v. Hence for all vertices v':

$$vR_{\pi(s,M)}v' \quad \text{iff} \quad vR_{\pi(s)\cup\pi(i_0)\cup\ldots\cup\pi(i_m);\pi(s,M)}v' \quad \text{iff} \quad vR_{\pi(s)\cup\pi(i_0)\cup\ldots\cup\pi(i_m)}\circ R_{\pi(s,M)}v'.$$

Now, either $P(v) \in M$ or $P(v) \notin M$. If the latter, for no $i \in M$ there is a v'' such that $vR_{\pi(i)}v''$. Hence, for an arbitrary vertex v'', we have $vR_{\pi(s)\cup\pi(i_0)\cup\ldots\cup\pi(i_m)}v''$ if and only if $vR_{\pi(s)}v''$. Consequently also, $vR_{\pi(s,M)}v''$ if and only if $vR_{\pi(s)} \circ R_{\pi(s,M)}v''$. Now consider the following equivalences:

$$\begin{array}{lll} vR_{\pi(s,M)}v' & \text{iff}_{P(v)\notin M} & vR_{\pi(s)} \circ R_{\pi(s,M)}v' \\ & \text{iff} & \text{for some } v''\text{: } vR_{\pi(s)}v''R_{\pi(s,M)}v' \\ & \text{iff}_{(*)} & s(v)\, R_{\pi(s,M)}v' \\ & \text{iff}_{i.h.} & v' \in \hat{s}_M(s(v)) \\ & \text{iff}_{(**)} & v' \in \hat{s}_M(v) \end{array}$$

The induction hypothesis is applicable because obviously $\mathit{hgt}(s(v)) < \mathit{hgt}(v)$. Observe further that in virtue of Definition 5.1, $vR_{\pi(s)}v''$ if and only if $v'' = s(v)$; whence the equivalence marked with the asterisk. The inference step indicated with the double asterisk is valid in virtue of $s_M(v) = \{s(v)\}$, because $P(v) \notin M$, and therefore $\hat{s}_M(v) = \bigcup\{\hat{s}_M(w) : \ w \in s_M(v)\} = \hat{s}_M(s(v))$.

In the former case in which $P(v) \in M$, let i denote $P(v)$. Then, also because $vR_{\pi(s)}v''$ implies $vR_{\pi(i)}v''$:

$$vR_{\pi(s)\cup\pi(i_0)\cup\ldots\cup\pi(i_m)}v'' \quad \text{iff} \quad vR_{\pi(s)\cup\pi(i)}v'' \quad \text{iff} \quad vR_{\pi(i)}v''.$$

Now consider the following equivalences:

$$\begin{array}{lll} vR_{\pi(s,M)}v' & \text{iff}_{P(v)=i} & vR_{\pi(i)} \circ R_{\pi(s,M)}v' \\ & \text{iff} & \text{for some } v''\text{: } vR_{\pi(i)}v''R_{\pi(s,M)}v' \\ & \text{iff}_{(*)} & \text{for some } v'' \in s_M(v) : v''R_{\pi(s,M)}v' \\ & \text{iff}_{i.h.} & \text{for some } v'' \in s_M(v) : v' \in \hat{s}_M(v'') \\ & \text{iff} & v' \in \bigcup\left\{\hat{s}_M\left(v''\right) : \ v'' \in s_M(v)\right\} \\ & \text{iff} & v' \in \hat{s}_M(v)\,. \end{array}$$

The induction hypothesis is applicable because for all $v'' \in s_M(v)$, it is the case that $\mathit{hgt}(v'') < \mathit{hgt}(v)$. Here, the inference step marked with the asterisk holds in virtue of Definition 5.1 and the definition of $s_M(v)$ on page 289, above. □

The construction of the frames $\mathfrak{F}_G^{\mathsf{MML}}$ and $\mathfrak{F}_G^{\mathsf{PDL}}$ from an extensive form G guarantees that if the one satisfies $(\hat{s}_i, i, \hat{s}_\varnothing)$-Euclidicity, the other satisfies satisfies $(\pi(s,i), i, \pi(s,\varnothing))$-Euclidicity and *vice versa*.

Corollary 5.1. Let G be an extensive game and consider the game-frame $\mathfrak{F}_G^{\mathsf{MML}}$ for a multi-modal matrix langauge $\mathcal{L}(A,B)$ and the dynamic game-frame $\mathfrak{F}_G^{\mathsf{PDL}}$ for a dynamic language $\mathcal{L}(A',B')$. Assume that $B_0 = B'_0$ and that $B_1 = B'_{11}$. Then for arbitrary $i \in B_0$ and $s \in B_1$:

$$\mathfrak{F}_G^{\mathsf{MML}} \text{ satisfies } (\hat{s}_i, i, \hat{s}_\varnothing)\text{-Euclidicity} \quad \text{iff} \quad \mathfrak{F}_G^{\mathsf{PDL}} \text{ satisfies } (\pi(s,i), i, \pi(s))\text{-Euclidicity.}$$

Proof:
Consider arbitrary vertices v and v' in the extensive game G. First observe that vR_iv' in $\mathfrak{F}_G^{\mathsf{MML}}$ if and only if vR_iv' in $\mathfrak{F}_G^{\mathsf{PDL}}$ by definition of game-frames and dynamic game-frames. Also for $X \subseteq \{i\}$:

$$vR_{\hat{s}_X}v' \text{ in } \mathfrak{F}_G^{\mathsf{MML}} \quad \text{iff}_{\text{Def. 3.1}} \quad v' \in \hat{s}_X(v) \quad \text{iff}_{\text{Prop. 5.1}} \quad vR_{\pi(s,X)}s \text{ in } \mathfrak{F}_G^{\mathsf{PDL}}.$$

Hence, in particular, $R_{\hat{s}_\emptyset} = R_{\pi(s,\emptyset)}$ and $R_{\hat{s}_i} = R_{\pi(s,i)}$. The claim then follows immediately. ◻

In virtue of this observation we now have the following result, which states that subgame perfect Nash equilibria can be characterized in dynamic multi-modal languages in much the same manner as that was the case for multi-modal matrix languages.

Corollary 5.2. Let $\mathcal{L}(A, B)$ be a dynamic multi-modal language and let $G \in \mathcal{G}_B$. Assume that B_0 contain a label for each player in G. Then, for s a strategy profile in G that is also a label in B_1 and M a subset of B_0 containing a label for each interested player of G:

$$s \text{ is a s.p. Nash equilibrium in } G \quad \text{iff} \quad \mathfrak{F}_G \Vdash \textstyle\bigwedge_{i\in M} \big(\langle\pi(s,i)\rangle[i]\varphi \to [\pi(s,\emptyset)]\varphi\big).$$

Proof:
Almost immediate from Theorem 3.1, Proposition 5.1 and the semantics of multi-modal matrix languages on game-frames. Consider an extensive game G as well as the following equivalences:

s is a s.p. Nash equilibrium	iff$_{\text{Th. 3.1}}$	$\mathfrak{F}_G^{\mathsf{MML}} \Vdash \bigwedge_{i\in M} \big(\langle\hat{s}_i\rangle[i]\varphi \to [\hat{s}_\emptyset]\varphi\big)$
	iff	for all $i \in M$: $\mathfrak{F}_G^{\mathsf{MML}} \Vdash \langle\hat{s}_i\rangle[i]\varphi \to [\hat{s}_\emptyset]\varphi$
	iff	for all $i \in M$, $\mathfrak{F}_G^{\mathsf{MML}}$ satisfies $(\hat{s}_i, i, \hat{s}_\emptyset)$-Euclidicity
	iff$_{\text{Coroll. 5.1}}$	for all $i \in M$, $\mathfrak{F}_G^{\mathsf{PDL}}$ satisfies $(\pi(s,i), i, \pi(s))$-Euclidicity
	iff	$\mathfrak{F}_G^{\mathsf{PDL}} \Vdash \bigwedge_{i\in M} \big(\langle\pi(s,i)\rangle[i]\varphi \to [\pi(s,\emptyset)]\varphi\big)$.

This concludes the proof. ◻

A dynamic game-frame of Definition 5.1 reflects the structure of the underlying extensive game in considerably finer detail than the game-frame of Definition 3.1 does for the same game. This feature, however, comes with a vengeance in that it imposes heavier requirements on the models to be constructed in a Henkin-style completeness proof. The issue as to a complete axiomatization of the dynamic framework with respect to dynamic game-frames we leave as an open question.

6. Conclusions and Other Topics

In this paper we proposed the use of multi-modal matrix languages for the formal description of a class of extensive games. The games in this particular class all had a finite horizon and assumed perfect information on the part of the players. By focussing on such a limited class of games, the correspondences between the games and the logic could be kept relatively simple. Independent issues were left out of the picture, so as to emphasize the fundamental idea of how modal languages can be used to describe extensive games. Thus, the analysis passed over fundamental game-theoretical topics such as coalition formation, mixed strategies and repeated and infinite games. Incorporation of these issues in the present framework warrants further investigation. Still, a proper treatment would quite likely demand considerable extensions of the languages presented in this paper. Special mention should be made of imperfect information and related epistemic issues, as analyses of knowledge and belief using modal logic have been firmly established within the field of Artificial Intelligence. Incorporation of epistemic logic in the present framework may lead to a more comprehensive analysis. The concomitant complications should not be shunned.

The multi-modal matrix languages were especially designed to deal with (subgame perfect) Nash equilibrium in pure strategies. Its expressive power is limited to preferences and *individual* divergences from a strategy profile. The characterization of other game-theoretical notions — such as *Pareto efficiency*, *dominance* as well as the various refinements of Nash equilibrium as they have been suggested in the literature — may require (slightly) more sophisticated concepts. More structure of the extensive games is preserved in the dynamic game-frames. Accordingly, we may expect more from the dynamic language of PDL as to expressiveness with respect to other game-theoretical concepts than Nash equilibrium alone.

These considerations put in perspective the multi-modal matrix languages as we proposed to use them in the description of extensive games. They should by no means be taken as a proposal for a comprehensive and ultimate logical language for the description of extensive games. Rather, we meant to expose some of the structural properties of extensive games which render some strategy profiles to be (subgame perfect) Nash equilibria. The fact that these properties are characterizable in quite an inelaborate formal language, says something fundamental about the elementary nature of Nash equilibria and the expressive requirements for a language to characterize them.

A. Completeness of Some Extensive Game Logics

This appendix concerns the completeness of F with respect to the class of models based on game-frames. On the basis of this proof we can also demonstrate completeness of the logics $\mathsf{F5}_{s,i}$ and $\mathsf{F5}_s^N$ with respect to the class of models on game-frames in which the strategy profile s contains a subgame perfect best-response for Player i and those in which s is a subgame perfect Nash-equilibrium. Before entering on the formal elaboration, we first devote some more or less informal remarks as to the structure of the proof.

The semantics for the multi-modal matrix languages is based on the notion of a game-frame. These game-frames constitute a proper subclass of the relational structures, also known as Kripke frames, multi-modal languages are commonly interpreted in. We prove completeness of the multi-modal logic F with respect to the class of game-frames as introduced in Definition 3.1:

$$\Gamma \models_{\mathsf{F}} \varphi \quad \text{implies} \quad \Gamma \vdash_{\mathsf{F}} \varphi.$$

By a standard argument it suffices to demonstrate that for every F-consistent theory there is a model on a game-frame that satisfies it. At page 298 we argued that it would be natural to adopt a method in which for each F-consistent theory separately a game-model $\mathfrak{M}_\Gamma^{\mathsf{F}}$ is constructed which satisfies Γ at the root. With a slight variation in the notion of consistency employed, this construction also provides a game-model for any other theory that is consistent in any other extensive game logic, as F is the smallest (or weakest) extensive game logic. For completeness of $\mathsf{F5}_{i,s}$ and $\mathsf{F5}_s^N$, however, it has additionally to be proved that the models this construction yields are in the appropriate class of models. *I.e.*, for $\mathsf{F5}_{i,s}$ it has to be shown that for each $\mathsf{F5}_{i,s}$-consistent theory Γ the model $\mathfrak{M}_{\Gamma}^{\mathsf{F5}_{i,s}}$ is based on an extensive game in which the strategy profile s contains a best response for player i. In the case of $\mathsf{F5}_s^N$, similarly, one should show that, in the extensive game underlying a model $\mathfrak{M}_{\Gamma}^{\mathsf{F5}_s^N}$, for each interested player for a $\mathsf{F5}_s^N$-consistent theory Γ, there be a label in N, and, moreover, the strategy profile denoted by the label s is a subgame perfect Nash equilibrium.

For each extensive game logic Λ in a multi-modal matrix language and each theory Γ we construct a model $\mathfrak{M}_{\Gamma}^{\Lambda}$, omitting the superscript Λ when clear from the context. Let Λ be an extensive game logic and Γ be a Λ-consistent theory in a multi-modal matrix language $\mathcal{L}(A, B)$. The main burden will be on guaranteeing that the model $\mathfrak{M}_{\Gamma}$ be an actual game model. For each Λ-consistent theory Γ, therefore, we construct a model $\mathfrak{M}_{\Gamma}$ along with an extensive game $\boldsymbol{G}_{\Gamma}$. A game-frame $\mathfrak{F}_{\boldsymbol{G}_{\Gamma}}$ based on this game $\boldsymbol{G}_{\Gamma}$ then underlies $\mathfrak{M}_{\Gamma}$. In an effort to avoid double subscripts, for each Λ-consistent theory Γ we will denote the frame $\mathfrak{F}_{\boldsymbol{G}_{\Gamma}}$ by $\mathfrak{F}_{\Gamma}$.

In the construction of the game-model $\mathfrak{M}_{\Gamma}$, we first define a labelled tree $\mathfrak{T}_{\Gamma}$ consisting of a tree $\boldsymbol{\Sigma}_{\Gamma}$ and a labelling function $\boldsymbol{\theta}_{\Gamma}$ assigning Λ-consistent theories to the vertices in $\boldsymbol{\Sigma}_{\Gamma}$. In particular, $\boldsymbol{\theta}_{\Gamma}$ assigns a maximal Λ-consistent extension of Γ to the root of $\boldsymbol{\Sigma}_{\Gamma}$. The set of vertices $\boldsymbol{\Sigma}_{\Gamma}$ is not entirely independent of the labelling function $\boldsymbol{\theta}_{\Gamma}$ since it may depend on the theory assigned to a particular vertex whether $\boldsymbol{\Sigma}_{\Gamma}$ should also contain another vertex. For this reason, induction is relied upon in the definition of $\mathfrak{T}_{\Gamma}$. This tree $\mathfrak{T}_{\Gamma}$ contains sufficient information for the definition of a fully-fledged extensive game denoted by $\boldsymbol{G}_{\Gamma}$ as well as that for the game model $\mathfrak{M}_{\Gamma}$ based on $\boldsymbol{G}_{\Gamma}$. The tree on which $\boldsymbol{G}_{\Gamma}$ is based is given by $\boldsymbol{\Sigma}_{\Gamma}$. The number of players in $\boldsymbol{G}_{\Gamma}$ turns out to be one greater than the number player labels in $\mathcal{L}(A, B)$. Also which player is to move at which node depends on the structure of $\boldsymbol{\Sigma}_{\Gamma}$. Finally, the players' preferences over the vertices of $\boldsymbol{\Sigma}_{\Gamma}$ are derived from the theories the labelling function $\boldsymbol{\theta}_{\Gamma}$ assign to the vertices of $\boldsymbol{\Sigma}_{\Gamma}$.

The vertices of the tree $\boldsymbol{\Sigma}_{\Gamma}$ are chosen in such a way that appropriate strategy profiles in $\boldsymbol{G}_{\Gamma}$ can easily be recovered to serve as the interpretations for the labels in B_1 of $\mathcal{L}(A, B)$. This makes that by Definition 3.1 the extensive game $\boldsymbol{G}_{\Gamma}$ defines unequivocally a *game-frame* $\mathfrak{F}_{\Gamma}$. A suitable valuation function I for $\mathfrak{F}_{\Gamma}$ is found by another appeal to the labelling function $\boldsymbol{\theta}_{\Gamma}$ of $\mathfrak{T}_{\Gamma}$: Let I map each vertex v of $\mathfrak{F}_{\Gamma}$ onto the set $\{a \in A : a \in \boldsymbol{\theta}_{\Gamma}(v)\}$. Then define $\mathfrak{M}_{\Gamma}$ as the very model on $\mathfrak{F}_{\Gamma}$ with I as valuation. The vertices of $\mathfrak{T}_{\Gamma}$ and $\mathfrak{M}_{\Gamma}$ coincide and we will prove that any formula in the theory associated with a vertex v in $\mathfrak{T}_{\Gamma}$, *i.e.*, the theory $\boldsymbol{\theta}_{\Gamma}(v)$, is satisfied at v in $\mathfrak{M}_{\Gamma}$. Because in $\mathfrak{T}_{\Gamma}$ the root is decorated with a maximal Λ-consistent extension of Γ, we may eventually conclude that $\mathfrak{M}_{\Gamma}$ satisfies Γ at the root node. The dependencies of the various elements of the construction of $\mathfrak{M}_{\Gamma}$— *viz.*, $\mathfrak{T}_{\Gamma}$, $\boldsymbol{G}_{\Gamma}$, $\mathfrak{F}_{\Gamma}$ and $\mathfrak{M}_{\Gamma}$ itself — are as depicted in Figure 7.

We begin the formal exposition by making some harmless but convenient assumptions. We assume to be working in a countable multi-modal matrix language $\mathcal{L}(A, B)$ with $B = N \cup S \cup (N \times S)$. Let further assume N to be given by a finite initial segment of the positive integers, *i.e.*, $N = \{1, \ldots, n\} \subseteq \omega$ with n the number of player-labels $\|N\|$. The game $\boldsymbol{G}_{\Gamma}$ to be constructed, will further turn out to comprise an additional *mystery player*, which will be denoted by 0. We will also assume an arbitrary but fixed enumeration $\varphi_0, \ldots, \varphi_n, \ldots$ of the formulas of $\mathcal{L}(A, B)$. Moreover, the concept of a maximal Λ-consistent extension of a theories will be heavily be relied upon. Here, the Lindenbaum lemma for Λ, stating that any Λ-consistent theory can be extended to a maximal Λ-consistent theory, is reproduced without the routine proof.

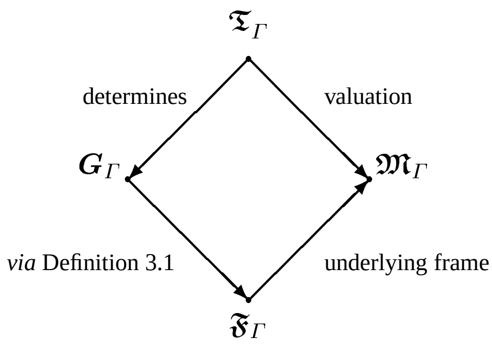

Figure 7. Structure of the construction of $\mathfrak{M}_\Gamma$, where $\mathfrak{T}_\Gamma = (\boldsymbol{\Sigma}_\Gamma, \boldsymbol{\theta}_\Gamma)$.

Fact A.1. (Lindenbaum lemma)
Every Λ-consistent theory in $\mathcal{L}(A, B)$ can be extended to a maximal Λ-consistent theory.

Having assumed a fixed enumeration of the formulas of $\mathcal{L}(A, B)$ we may for each extensive game logic Λ assume a *closure operator* Cl_Λ mapping each Λ-consistent theory Γ onto a *unique* maximal Λ-consistent extension of Γ. Usually, we will omit the subscript, when the logic Λ is understood from the context. The construction of the models $\mathfrak{M}_\Gamma^\Lambda$ is uniform for all extensive game logics Λ *modulo* the notion of consistency involved.

The vertices of the labelled tree $\mathfrak{T}_\Gamma$ will be drawn from a set of finite strings over a particular set. Suppressing the implicit ordering of the strings we will deliberately confuse the tree $\boldsymbol{\Sigma}_\Gamma$ and its set of vertices. For each natural number $n \in \omega$, we let x^n denote the string of n occurrences of x, *e.g.*, $x^3 = xxx$ and $yx^2z = yxxz$. Let furthermore $|\sigma|$ denote the *length* of a string σ. We have ϵ stand for the empty string. In the remainder strings are assumed to be ordered by the *immediate prefix relation* $\prec$, defined for strings σ and σ' over a set X in such a way that $\sigma \prec \sigma'$ if and only if there is some x in X with $\sigma x = \sigma'$. *E.g.*, the strings xy and xyz are thus related but yy and yxx are not.

The vertices of the labelled tree $\mathfrak{T}_\Gamma$ are selected from the set of finite sequences over $T \cup S \cup \omega$ and includes the empty sequence ϵ. The sets T, S and ω are assumed to be pairwise disjoint, T to be countably infinite, and S the set of strategy-labels of $\mathcal{L}(A, B)$. We will assume there to be a partition of T in a countably infinite number of countably infinite blocks. Hence, for each sequence σ in $(T \cup S \cup \omega)^*$ we may assume there to be a unique countably infinite subset of T, denoted by T_σ and enumerated as $t_0^\sigma, \ldots, t_n^\sigma, \ldots$.

Conceptually, in the game $\boldsymbol{G}_\Gamma$ to be constructed, the elements of $T \cup S \cup \omega$ could be taken as possible actions and each sequence σ is to represent a *history of play.* A sequence $ts's$ is then the vertex that will be reached if subsequently the 'actions' t, s' and s are being played. In the game $\boldsymbol{G}_\Gamma$ to be defined the strategy profile the label s is to represent can then easily be recovered as the function that maps each sequence σ onto σs. The sets T and ω are added in order to assure that in $\boldsymbol{G}_\Gamma$ there be a sufficient number vertices, and *a fortiori* a sufficient number of strategy profiles. Roughly speaking, elements of T are used to introduce vertices falsifying φ as witnesses for formulas of the form $\neg[i]\varphi$. Similarly, the elements of ω are used to construct witness states for formulas of the form $\neg[\hat{s}_i]\varphi$.

Now the stage has been set for the definition of the labelled tree $\boldsymbol{\mathfrak{T}}_\Gamma$. The fundamental idea is that the for each formula of the form $\neg[\beta]\varphi$ in a theory associated with a vertex σ by $\boldsymbol{\theta}_\Gamma$, there should also be a vertex σ' that falsifies φ and, moreover, is reachable by the accessibility relation R_β in the model $\mathfrak{M}_\Gamma$. The theory associated with vertex σ' will thus contain $\neg\varphi$ and will also have to comply to certain consistency constraints as is usual in a construction method. As, for different labels β and β' the theories $\{\psi : \ [\beta]\psi \in \Gamma\}$ and $\{\psi : \ [\beta']\psi \in \Gamma\}$ are not in general Λ-compatible, proper care should be taken that no leaf be reachable by two different accessibility relations R_β and $R_{\beta'}$ from a vertex such that $\boldsymbol{\theta}_\Gamma(v)$ may contain both formulas of the form $[\beta]\psi$ and $[\beta']\psi$. Simultaneously, it has to be ascertained that the tree constructed is a game-tree in accordance with the interpretation of strings as actions and the strategy profiles as described above. Thus, for each internal node σ of $\boldsymbol{\mathfrak{T}}_\Gamma$ and each label in B_1 of $\mathcal{L}(A, B)$, there should be a unique leaf σs^n, for some appropriate integer $n \in \omega$, representing the outcome of the strategy profile of $\boldsymbol{G}_\Gamma$ the label s stands for in the model $\mathfrak{M}_\Gamma$. Moreover, care should be taken that the tree $\boldsymbol{\mathfrak{T}}_\Gamma$ have a finite horizon.

For each t in $\{\epsilon\} \cup T$ and each theory Γ in $\mathcal{L}(A, B)$, we first define inductively for each $n \in \omega$ a tree $\mathfrak{T}^n_{\Gamma,t}$ (no boldface!) the vertices of which are decorated with theories in $\mathcal{L}(A, B)$. Then the tree $\mathfrak{T}_{\Gamma,t}$ is defined as the limit of this induction. This latter kind of labelled tree will form the modules which eventually compose the tree $\boldsymbol{\mathfrak{T}}_\Gamma$ for Γ. The set of vertices of $\mathfrak{T}_{\Gamma,t}$ we denote by $\Sigma_{\Gamma,t}$ and the labelling function assigning theories of $\mathcal{L}(A, B)$ to the vertices in $\Sigma_{\Gamma,t}$ by $\theta_{\Gamma,t}$. The root of $\mathfrak{T}_{\Gamma,t}$ is taken to be t and the other vertices and the theories assigned to them are chosen in accordance with the idea that Γ be eventually satisfiable at t.

At the basis of the induction the tree $\mathfrak{T}^0_{\Gamma,t}$ is defined, with $\Sigma^0_{\Gamma,t}$ as vertices and $\theta^0_{\Gamma,t}$ as labelling function. The idea is that $\mathfrak{T}^0_{\Gamma,t}$ contain, for each s in S, a *unique* leaf that can be taken as the outcome of s when play is commenced in the root t. The design is such that along any such path each of the players in N is to move once. In general, player i is assumed to move at ts^i. In particular, the mystery player 0 thus makes a decision at the root t. Hence, $\Sigma^0_{\Gamma,t}$ contains t as well as each sequence ts^n with $n \leqslant \|N\| + 1$, the idea being that each player $i \in N$ is to move at ts^i and $ts^{\|N\|+1}$ is the outcome s determines in t. By definition, a strategy profile should be defined at each internal node. Hence for each ts^i with $i \leqslant \|N\|$ and each label s' in S different from s we also distinguish a leaf $ts^i s'$ in $\Sigma^0_{\Gamma,t}$. No further vertices are in $\Sigma^0_{\Gamma,t}$. Figure 8 depicts the tree for $\Sigma^0_{\Gamma,t}$ in a language containing two labels for players and also two labels for (the outcome functions of) strategy profiles.

The labelling function $\theta^0_{\Gamma,t}$ assigns the maximal Λ-consistent extension $Cl(\Gamma)$ of Γ to the root t and the empty theory to each of the internal vertices. However, arbitrary the latter may seem, it will prove to be convenient as the proof develops. If now the theory $Cl(\Gamma)$ is supposed to be satisfied at t in some game-model on $\Sigma^0_{\Gamma,t}$, then each leaf $ts^{\|N\|+1}$ should satisfy a formula ψ whenever $[\hat{s}_\emptyset]\psi$ is in $Cl(\Gamma)$. Similarly, if $Cl(\Gamma)$ contains a formula $[\hat{s}_i]\psi$, then each leaf $ts^i s'$ should satisfy ψ. The assignment function $\theta^0_{\Gamma,t}$ is fixed accordingly.

The role of the mystery player can now also be clarified. Suppose t had been assigned to a player i with a label in N but different from 1. Then, for a strategy profiles s and s' with a labels in B_1, the leaf tss' would eventually be reachable not only from t by $R_{\hat{s}_1}$, as intended, but also by $R_{\hat{s}'_i}$! If so, the theory associated with tss' should then contain all formulas ψ for which either $[\hat{s}_1]\psi$ or $[\hat{s}'_i]\psi$. The Λ-consistency of a such a theory, however, cannot in general be guaranteed. By assigning a player without a label in N to the root t this contingency will not occur.

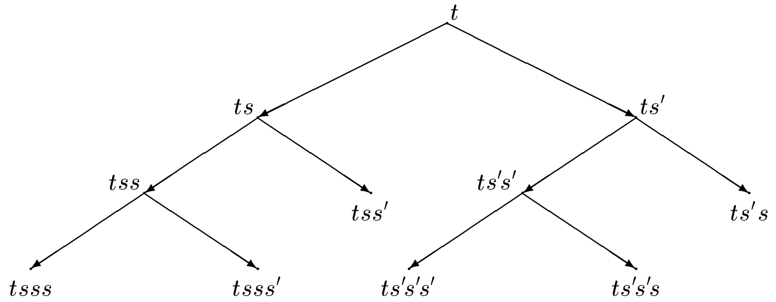

Figure 8. The tree $\Sigma^0_{\Gamma,t}$ for language $\mathcal{L}(A,B)$ with $N=\{1,2\}$ and $S=\{s,s'\}$.

Formally we define:

$$\mathfrak{T}^0_{\Gamma,t} \quad =_{df.} \quad \left(\Sigma^0_{\Gamma,t}, \theta^0_{\Gamma,t}\right),$$

where $\Sigma^0_{\Gamma,t}$ is a subset of $(T \cup S \cup \omega)^*$ and $\theta^0_{\Gamma,t}$ a function assigning theories in $\mathcal{L}(A,B)$ to the sequences σ in $\Sigma^0_{\Gamma,t}$:

$$\Sigma^0_{\Gamma,t} \quad =_{df.} \quad \{t\} \cup \{ts^i s' : \ s, s' \in S \ \text{ and } \ 0 \leqslant i \leqslant \|N\| \}$$

$$\theta^0_{\Gamma,t}(\sigma) \quad =_{df.} \quad \begin{cases} Cl(\Gamma) & \text{if } \sigma = t \\ Cl\left(\{\psi : \ [\hat{s}_\emptyset]\psi \in Cl(\Gamma)\}\right) & \text{if } \sigma = ts^{\|N\|+1} \\ Cl\left(\{\psi : \ [\hat{s}_i]\psi \in Cl(\Gamma)\}\right) & \text{if } \sigma = ts^i s' \text{ and } s \neq s' \\ \emptyset & \text{otherwise.} \end{cases}$$

In the inductive step, defining $\mathfrak{T}^{n+1}_{\Gamma,t}$ from $\mathfrak{T}^n_{\Gamma,t}$, we check whether φ_n, the $n-1$-st formula in the enumeration, is of the form $\neg[\hat{s}_i]\psi$. If it is and φ_n moreover occurs in $Cl(\Gamma)$, the string $ts^i n$ is added to the set of vertices and assigned the maximal Λ-extension of $\{\neg\psi\} \cup \{\chi : \ [\hat{s}_i]\chi \in \Gamma\}$. In any other case $\mathfrak{T}^{n+1}_{\Gamma,t}$ and $\mathfrak{T}^n_{\Gamma,t}$ are identical.

The reason for this is that if the root t were to satisfy a formula of the form $\neg[\hat{s}_i]\psi$, then a leaf v' should be reachable from t *via* $R_{\hat{s}_i}$ and not satisfy ψ. Being reachable *via* $R_{\hat{s}_i}$, the leaf v' should, additionally, also satisfy any formula χ that is such that t forces $[\hat{s}_i]\chi$. Note that a similar construction is unnecessary if φ_n is of the form $\neg[\hat{s}_\emptyset]\psi$. In virtue of Axiom $D!_{\hat{s}_\emptyset}$, the formula $\neg[\hat{s}_\emptyset]\psi$ is equivalent to $[\hat{s}_\emptyset]\neg\psi$ and the latter is thus element of $Cl(\Gamma)$. Hence, this case has already been taken care of by the construction of $\mathfrak{T}^0_{\Gamma,t}$. A similar remark applies to the case in which a formula of the form $\neg[\hat{s}_\emptyset]\psi$ or $\neg[\hat{s}_i]\psi$ is contained in the theory assigned to a leaf σ in $\Sigma^n_{\Gamma,t}$ by $\theta^n_{\Gamma,t}$. Then, Axiom $F5_{\hat{s}_X,\hat{s}_Y}$ makes that already $\neg\psi$ is in $\theta^n_{\Gamma,t}(\sigma)$. The internal vertices contain no formulas, let alone formulas that require "*witness*" states.

Formally define:

$$\mathfrak{T}^{n+1}_{\Gamma,t} \quad =_{df.} \quad \left(\Sigma^{n+1}_{\Gamma,t}, \theta^{n+1}_{\Gamma,t}\right),$$

where:

$$\Sigma^{n+1}_{\Gamma,t} =_{df.} \begin{cases} \Sigma^n_{\Gamma,t} \cup \{ts^i n\} & \text{if } \varphi_n = \neg[\hat{s}_i]\psi \text{ and } \varphi_n \in Cl(\Gamma) \\ \Sigma^n_{\Gamma,t} & \text{otherwise.} \end{cases}$$

$$\theta^{n+1}_{\Gamma,t}(\sigma) =_{df.} \begin{cases} Cl\left(\{\neg\psi\} \cup \{\chi : [\hat{s}_i]\chi \in Cl(\Gamma)\}\right) & \text{if } \sigma = ts^i n, \varphi_n = \neg[\hat{s}_i]\psi \text{ and } \varphi_n \in Cl(\Gamma) \\ \theta^n_{\Gamma,t}(\sigma) & \text{otherwise, } \textit{i.e.}, \text{ if } \sigma \in \Sigma^n_{\Gamma,t}. \end{cases}$$

Finally define $\mathfrak{T}_{\Gamma,t}$ as:

$$\mathfrak{T}_{\Gamma,t} =_{df.} \Big(\bigcup_{n\in\omega} \Sigma^n_{\Gamma,t}, \bigcup_{n\in\omega} \theta^n_{\Gamma,t}\Big).$$

The theories $\theta_{\Gamma,t}$ assigns to the vertices may also contain formulas of the form $\neg[i]\psi$. If this is the case for a vertex σ, then the construction should also contain a vertex associated with a maximal Λ-consistent theory containing $\neg\psi$ as well as any formula χ if $[i]\chi$ is in the theory associated with σ. To accommodate this type of formula, we push the construction one step further.

We now define inductively for each $n \in \omega$ a *collection of decorated trees of the form* $\mathfrak{T}_{\Gamma,t}$, from which we eventually manufacture the tree $\boldsymbol{\mathfrak{T}}_\Gamma$. At the basis, this collection consists of the tree $\mathfrak{T}_{\Gamma,\epsilon}$, which has Γ itself associated with its root ϵ. The empty sequence ϵ being a prefix to any sequence, ϵ will also be the root of the tree to be constructed and eventually also of the game-model $\mathfrak{M}_\Gamma$. If any of the vertices σ of $\mathfrak{T}_{\Gamma,\epsilon}$ contains a formula φ_k of the form $\neg[i]\psi$ a new tree $\mathfrak{T}_{\Theta,t^\sigma_k}$ is added to the collection, on the understanding that the theory Θ equals $\{\neg\psi\} \cup \{\chi : [i]\chi \in \theta(\sigma)\}$. Then this process is repeated for the new collection *ad infinitum*. Thus define:

$$\boldsymbol{\mathfrak{T}}^0_\Gamma =_{df.} \{\mathfrak{T}_{\Gamma,\epsilon}\}$$

$$\boldsymbol{\mathfrak{T}}^{n+1}_\Gamma =_{df.} \boldsymbol{\mathfrak{T}}^n_\Gamma \cup \bigcup_{(\Sigma,\theta)\in\boldsymbol{\mathfrak{T}}^n_\Gamma} \Big\{\mathfrak{T}_{\Theta,t^\sigma_k} : \sigma \in \Sigma \text{ and } \varphi_k = \neg[i]\psi \text{ and } \varphi_k \in \theta(\sigma)\Big\}$$

$$\text{where,} \quad \Theta = \{\neg\psi\} \cup \{\chi : [i]\chi \in \theta(\sigma)\}.$$

We will assume that T has been chosen and distributed over the sequences in $(T \cup S \cup \omega)^*$ in such a way that for each n the t^σ_k-s in the above definition are "*fresh*". This guarantees that at each stage of the induction the set of vertices of each tree added to the collection is disjoint from any other set of vertices of a tree thus introduced, as well as from any set of vertices of a tree already present in the collection. Hence, the domains of the various functions θ remain separate as well. Then set:

$$\boldsymbol{\mathfrak{T}}^\omega_\Gamma =_{df.} \bigcup_{n\in\omega} \boldsymbol{\mathfrak{T}}^n_\Gamma.$$

We are now in a position to formally define the decorated tree $\boldsymbol{\mathfrak{T}}_\Gamma$ as:

$$\boldsymbol{\mathfrak{T}}_\Gamma =_{df.} (\boldsymbol{\Sigma}_\Gamma, \boldsymbol{\theta}_\Gamma),$$

where:

$$\boldsymbol{\Sigma}_\Gamma =_{df.} \bigcup \{\Sigma : (\Sigma,\theta) \in \boldsymbol{\mathfrak{T}}^\omega_\Gamma\}$$

$$\boldsymbol{\theta}_\Gamma =_{df.} \bigcup \{\theta : (\Sigma,\theta) \in \boldsymbol{\mathfrak{T}}^\omega_\Gamma\}.$$

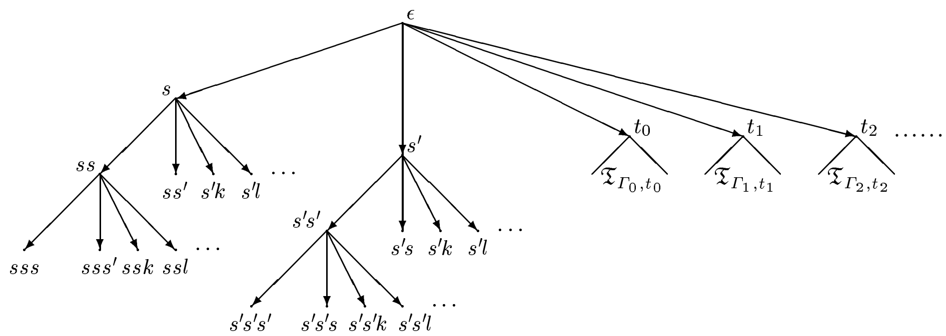

Figure 9. The tree $\mathfrak{T}_\Gamma$ for a theory Γ in a language $\mathcal{L}(A,B)$ with $N = \{1,2\}$ and $S = \{s, s'\}$. Each subtree $\mathfrak{T}_{\Gamma_k,t_k}$ is introduced in virtue of a formula of the form $\neg[i]\psi$ being in the theory assigned to a vertex in another part of the tree.

Before we proceed, the reader be assured that, for Γ a Λ-consistent theory, $\boldsymbol{\mathfrak{T}}_\Gamma$ is indeed a tree with its vertices $\boldsymbol{\Sigma}_\Gamma$ ordered by the immediate prefix relation $\prec$ and that $\boldsymbol{\theta}_\Gamma$ is a *function* assigning a unique theory in $\mathcal{L}(A,B)$ to each vertex σ in $\boldsymbol{\Sigma}_\Gamma$.

Fact A.2. For each Λ-consistent theory Γ, the set $\boldsymbol{\Sigma}_\Gamma$ is a tree if ordered by the immediate prefix relation $\prec$, and $\boldsymbol{\theta}_\Gamma$ is a total function on $\boldsymbol{\Sigma}_\Gamma$.

Proof:
(*Sketch.*) For each $x \in \{\epsilon\} \cup T$, $\Sigma^0_{\Gamma,t}$ is defined as a tree ordered by $\prec$. Then observe that $\Sigma^{n+1}_{\Gamma,t}$ is obtained from $\Sigma^n_{\Gamma,t}$ by adding at most a fresh vertex ts^in. With ts^i already contained in $\Sigma^0_{\Gamma,t}$, we may by the induction hypothesis conclude that $\Sigma^{n+1}_{\Gamma,t}$ is a tree ordered by $\prec$ as well. Now consider $\boldsymbol{\mathfrak{T}}^\omega_\Gamma$. The "*freshness*" assumptions for the t^σ_n guarantee that for any two (Σ,θ) and (Σ',θ') in $\boldsymbol{\mathfrak{T}}^\omega_\Gamma$, the sets Σ and Σ' are disjoint. Finally, recognize that each $\mathfrak{T}_{\Theta,t}$ has t as root and that ϵ is an immediate prefix of t. This makes that $\boldsymbol{\Sigma}_\Gamma$ is a tree ordered by $\prec$. By a similar argument it can be shown that θ is functional for each (Σ,θ) in $\boldsymbol{\mathfrak{T}}^\omega_\Gamma$. The functionality of $\boldsymbol{\theta}_\Gamma = \bigcup\{\theta : (\Sigma,\theta) \in \boldsymbol{\mathfrak{T}}^\omega_\Gamma\}$ is guaranteed by the "*freshness*" assumptions for the t^σ_n. $\square$

Figure 9 depicts the structure of $\boldsymbol{\mathfrak{T}}_\Gamma$ for a language $\mathcal{L}(A,B)$ with $N = \{1,2\}$ and $S = \{s,s'\}$.

For technical convenience we will distinguish particular subsets of vertices in $\boldsymbol{\mathfrak{T}}_\Gamma$. First, the root node ϵ and the vertices t^σ_n that are the roots of subtrees $\mathfrak{T}_{\Gamma,t^\sigma_n}$ in $\boldsymbol{\mathfrak{T}}_\Gamma$ are collected in $\boldsymbol{T}_\Gamma$, *i.e.*,

$$\boldsymbol{T}_\Gamma \;=_{\textit{df.}}\; \{\epsilon\} \cup (T \cap \boldsymbol{\Sigma}_\Gamma)\,.$$

For each t in $\boldsymbol{T}_\Gamma$ the set of internal vertices and the set of leaves of the respective subtree $\mathfrak{T}_{\Theta,t}$ of $\boldsymbol{\mathfrak{T}}_\Gamma$ are denoted by $\boldsymbol{I}^t_\Gamma$ and $\boldsymbol{L}^t_\Gamma$, respectively. Obviously, $\boldsymbol{I}^t_\Gamma$ and $\boldsymbol{L}^t_\Gamma$ are disjoint and together exhaust the vertices in $\mathfrak{T}_{\Gamma',t}$. Finally, $\boldsymbol{W}^t_\Gamma$ comprises t together with the leaves $\boldsymbol{L}^t_\Gamma$. Some reflection reveals that $\boldsymbol{W}^t_\Gamma$

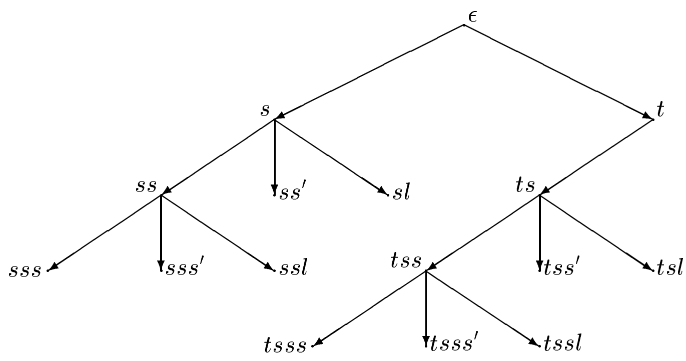

Figure 10. Types of vertex in $\mathfrak{T}_\Gamma$. The set $\boldsymbol{T}$ contains ϵ and (vertices of type) t. The vertices ϵ, s and ss make up $\boldsymbol{I}^{\epsilon}$, whereas t, ts and tss are the sole elements in $\boldsymbol{I}^{t}$. $\boldsymbol{L}^t$ is given by $\{tss', tsl, tsss, tsss', tssl\}$ and $\boldsymbol{W}$ contains all vertices except s, ss, ts and tss.

are exactly those vertices in $\mathfrak{T}_{\Gamma',t}$ which are labelled with a maximal consistent theory.

$$\boldsymbol{I}^t_\Gamma \;=_{df.}\; \left\{ts^i \in \boldsymbol{\Sigma}_\Gamma :\; s \in S \text{ and } i \leqslant \|N\|\right\}$$
$$\boldsymbol{L}^t_\Gamma \;=_{df.}\; \left\{ts^{\|N\|+1}, ts^i s', ts^i l \in \boldsymbol{\Sigma}_\Gamma :\; s \neq s' \text{ and } l \in \omega\right\}$$
$$\boldsymbol{W}^t_\Gamma \;=_{df.}\; \{t\} \cup \boldsymbol{L}^t_\Gamma$$

Let further $\boldsymbol{I}_\Gamma =_{df.} \bigcup_{t\in\boldsymbol{T}_\Gamma} \boldsymbol{I}^t_\Gamma$, $\boldsymbol{L}_\Gamma =_{df.} \bigcup_{t\in\boldsymbol{T}_\Gamma} \boldsymbol{L}^t_\Gamma$, and $\boldsymbol{W}_\Gamma =_{df.} \bigcup_{t\in\boldsymbol{T}_\Gamma} \boldsymbol{W}^t_\Gamma$. Since the vertices of all the trees in $\mathfrak{T}^\omega_\Gamma$ are pairwise disjoint $\boldsymbol{W}_\Gamma$ contains precisely those vertices labelled with maximal consistent theories. Moreover, $\boldsymbol{L}_\Gamma$ collects all leaves of $\mathfrak{T}_\Gamma$ and $\boldsymbol{I}_\Gamma$ its internal vertices and as such they are disjunct. Finally, we claim without giving a proof that:

$$\boldsymbol{\Sigma}_\Gamma \;=\; \boldsymbol{I}_\Gamma \cup \boldsymbol{L}_\Gamma.$$

The various "*types*" of sequence the collection $\boldsymbol{\Sigma}_\Gamma$ contains and how they relate is illustrated in Figure 10, for the same language as that of Figure 9.

Fact A.3. Let Γ be a Λ-consistent theory in $\mathcal{L}(A, B)$. For each $\sigma \in \boldsymbol{\Sigma}_\Gamma$, then, $\boldsymbol{\theta}_\Gamma(\sigma)$ is a maximal Λ-consistent theory if $\sigma \in \boldsymbol{W}_\Gamma$, and the empty theory, otherwise.

Proof:
(*Sketch.*) First we observe that for any extensive game logic Λ it is the case that:

$$\Gamma \nvdash \bot \quad \text{implies} \quad \{\psi :\; [\hat{s}_i]\psi \in \Gamma\} \nvdash \bot,$$
$$\Gamma \nvdash \bot \quad \text{implies} \quad \{\psi :\; [\hat{s}_\emptyset]\psi \in \Gamma\} \nvdash \bot,$$
$$\Gamma \cup \{\neg[i]\chi\} \nvdash \bot \quad \text{implies} \quad \{\neg\chi\} \cup \{\psi :\; [i]\psi \in \Gamma\} \nvdash \bot,$$
$$\Gamma \cup \{\neg[\hat{s}_i]\chi\} \nvdash \bot \quad \text{implies} \quad \{\neg\chi\} \cup \{\psi :\; [\hat{s}_i]\psi \in \Gamma\} \nvdash \bot.$$

The logic Λ being a normal logic, the latter two can be proved by a standard argument (*cf.*, [4], pp.198-199). The first and the second item are proved by an analogous argument, though essentially involving the axioms $D!_{\beta_1}$ and $F1_{(\beta_0,\beta_1),\beta_1}$. Again assume the contrapositive that $\{\psi : [\hat{s}_i]\psi \in \Gamma\} \vdash \bot$. Then there is a finite number of formulas $\psi_0, \dots, \psi_n \in \{\psi : [\hat{s}_i]\psi \in \Gamma\}$ such that $\psi_0, \dots, \psi_n \vdash \bot$. Consider the following implications:

$$\psi_0, \dots, \psi_n \vdash \bot \quad \text{implies} \quad [\hat{s}_i]\psi_0, \dots, [\hat{s}_i]\psi_n \vdash [\hat{s}_i]\bot \quad \text{implies} \quad \Gamma \vdash [\hat{s}_i]\bot$$
$$\text{implies}_{\,F1_{\hat{s}_i,\hat{s}_\emptyset}} \quad \Gamma \vdash [\hat{s}_\emptyset]\bot \quad \text{implies}_{\,D!_{\hat{s}_\emptyset}} \quad \Gamma \vdash \langle\hat{s}_\emptyset\rangle\bot \quad \text{implies} \quad \Gamma \vdash \bot.$$

The first and last implication are in virtue of Λ being a normal modal logic. The second item can be obtained using *Nec.* for $[\hat{s}_\emptyset]$ instead of for $[\hat{s}_i]$ and omitting the application of $F1_{\hat{s}_i,\hat{s}_\emptyset}$.

An inductive check of the construction of the various $\mathfrak{T}_{\Gamma,t}$ and eventually of that of $\boldsymbol{\mathfrak{T}}_\Gamma$, here omitted, then establishes the fact. □

For each leaf σ in $\boldsymbol{L}_\Gamma$ and $x \in \boldsymbol{T}_\Gamma$, we define a label β_σ in B of the language $\mathcal{L}(A, B)$, as follows:

$$\beta_\sigma \quad =_{df.} \begin{cases} \hat{s}_i & \text{if } \sigma = xs^is', \text{ for some } s' \neq s, \\ \hat{s}_i & \text{if } \sigma = xs^il, \text{ for some } l \in \omega \text{ and } i \leqslant \|N\|, \\ \hat{s}_\emptyset & \text{otherwise, } i.e., \text{ if } \sigma = xs^{\|N\|+1}. \end{cases}$$

The purpose of this definition is the following useful fact; its proof amounts to an easy check and is duly omitted.

Fact A.4. Let $x \in \boldsymbol{T}_\Gamma$. Then for all leaves $\sigma \in \boldsymbol{L}_\Gamma^x$: $\quad \{\varphi : [\beta_\sigma]\varphi \in \boldsymbol{\theta}_\Gamma(x)\} \subseteq \boldsymbol{\theta}_\Gamma(\sigma)$.

The intuition behind this definition is that in the game $\boldsymbol{G}_\Gamma$ to be constructed on basis of the labelled tree $\boldsymbol{\mathfrak{T}}_\Gamma$, for each *fixed* label s in S, the set of leaves in $\{ts^{\|N\|+1}, ts^is', ts^il : \in \boldsymbol{\Sigma}_\Gamma : s \neq s' \text{ and } l \in \omega\}$ will coincide with the leaves that can be reached from $x \in \boldsymbol{T}_\Gamma$ if at most one player deviates from the corresponding strategy profile s. Similarly, $s^{\|N\|+1}$ will be the outcome s determines at x. *I.e.*, in symbols, $\hat{s}_i(x) = \{ts^{\|N\|+1}, ts^is', ts^il : s \neq s' \text{ and } l \in \omega\}$ and $\hat{s}_\emptyset(x) = \{xs^{\|N\|+1}\}$. Since $\hat{s}_i$ and $\hat{s}_\emptyset$ in turn determine the accessibility relations for the corresponding labels in the game-frame on $\boldsymbol{G}_\Gamma$, Fact A.4 partly guarantees that for each sequence σ in $\boldsymbol{L}_\Gamma$ the theories $\boldsymbol{\theta}_\Gamma$ assigns to sequences are included in the theories that those sequences satisfy in a particular game-model on $\boldsymbol{G}_\Gamma$.

We are now in a position to define for each Λ-consistent theory Γ an extensive game $\boldsymbol{G}_\Gamma$ on basis of the labelled tree $\boldsymbol{\mathfrak{T}}_\Gamma$.

Definition A.1. (Extensive games for Λ-consistent theories)
Let Λ be an extensive game logic in a multi-modal matrix language $\mathcal{L}(A, B)$. Let, furthermore, Γ be a maximal Λ-consistent theory. We define: the extensive game $\boldsymbol{G}_\Gamma^\Lambda$ as:

$$\boldsymbol{G}_\Gamma^\Lambda \quad =_{df.} \quad (\boldsymbol{V}_\Gamma, \boldsymbol{R}_\Gamma, \boldsymbol{N}_\Gamma, \boldsymbol{P}_\Gamma, \{\boldsymbol{\rho}_i\}_{i\in\boldsymbol{N}_\Gamma}),$$

where $\boldsymbol{V}_\Gamma$ is the set of sequences $\boldsymbol{\Sigma}_\Gamma$ in $\boldsymbol{\mathfrak{T}}_\Gamma$ as above and $\boldsymbol{R}_\Gamma$ is given by the immediate prefix relation $\prec$ on $\boldsymbol{\Sigma}_\Gamma$, *i.e.*: for all $\sigma, \sigma' \in \boldsymbol{V}_\Gamma$

$$\sigma \boldsymbol{R}_\Gamma \sigma' \quad \text{iff} \quad \text{for some } x \in S \cup T \cup \omega : \quad \sigma' = \sigma x.$$

The players of the game $\boldsymbol{G}_\Gamma$ are given by the labels in N plus a *mystery player* 0, *i.e.*,

$$\boldsymbol{N}_\Gamma \quad =_{\mathit{df.}} \quad N \cup \{0\}\,.$$

The player assignment function $\boldsymbol{P}_\Gamma$ is such that for each internal vertex σ of $\boldsymbol{V}_\Gamma$, *i.e.*, $\sigma \in \boldsymbol{I}_\Gamma$:

$$\boldsymbol{P}_\Gamma(\sigma) = i \quad \text{iff} \quad \sigma = xs^i, \text{for some } x \in \boldsymbol{T}_\Gamma$$

Finally, the preferences of each player i in N are such that for all vertices $\sigma, \sigma' \in \boldsymbol{V}_\Gamma$:

$$\left(\sigma, \sigma'\right) \in \boldsymbol{\rho}_i \quad \text{iff} \quad \text{for all formulas } \varphi\text{:} \quad [i]\varphi \in \boldsymbol{\theta}_\Gamma(\sigma) \;\; \text{implies} \;\; \varphi \in \boldsymbol{\theta}_\Gamma\left(\sigma'\right)$$

The mystery player 0 we assume to be entirely indifferent, *i.e.*, $\boldsymbol{\rho}_0 =_{\mathit{df.}} \boldsymbol{V}_\Gamma \times \boldsymbol{V}_\Gamma$. When no confusion is likely, we will omit the subscript Γ as well as the superscript Λ in $\boldsymbol{G}_\Gamma^\Lambda$.

Fact A.5. For each maximal Λ-consistent theory Γ in $\mathcal{L}(A, B)$, $\boldsymbol{G}_\Gamma$ is a properly defined extensive game.

Proof:
Consider an arbitrary theory Γ in $\mathcal{L}(A, B)$. By construction $(\boldsymbol{V}, \boldsymbol{R})$ is a tree and some reflection reveals that the length of a string in $\boldsymbol{V}$ is no longer than $\|N\| + 2$. Hence, with N being finite, $(\boldsymbol{V}, \boldsymbol{R})$ has a finite horizon. By the same token the set of players $\boldsymbol{N}$ of $\boldsymbol{G}_\Gamma$ is finite as well, since $\boldsymbol{N}$ contains just one element more than N, *viz.*, the mystery player. The player assignment function $\boldsymbol{P}$ can readily be ascertained to be total on the internal vertices.

Finally, each of the players' preferences are reflexive, transitive and connected as required. For the mystery player 0 this is immediate, its preference relation being the universal relation over $\boldsymbol{V}$. So consider, for the remainder of the proof, an arbitrary player i in N. For reflexivity, observe that for any string $\sigma \in \boldsymbol{V}$ such that $\boldsymbol{\theta}(\sigma)$ is empty, $(\sigma, \sigma') \in \boldsymbol{\rho}_i$ is trivial. Otherwise, by Fact A.3, the theory $\boldsymbol{\theta}(\sigma)$ is maximal Λ-consistent. Then, assume for an arbitrary formula φ that $[i]\varphi \in \boldsymbol{\theta}(\sigma)$. With the axiom T_i then also $\varphi \in \boldsymbol{\theta}(\sigma)$. Consequently, $(\sigma, \sigma) \in \boldsymbol{\rho}_i$. For transitivity a similar run-of-the-mill argument suffices. Assume $(\sigma, \sigma'), (\sigma', \sigma'') \in \boldsymbol{\rho}_i$. Again if $\boldsymbol{\theta}(\sigma)$ is empty, immediately $(\sigma, \sigma'') \in \boldsymbol{\rho}_i$, as well. Otherwise, $\boldsymbol{\theta}(\sigma)$ is maximal Λ-consistent. Assume for an arbitrary formula φ that $[i]\varphi \in \boldsymbol{\theta}(\sigma)$. By axiom 4_i, also $[i][i]\varphi \in \boldsymbol{\theta}(\sigma)$. By definition of $\boldsymbol{\rho}_i$, then subsequently $[i]\varphi \in \boldsymbol{\theta}(\sigma')$ and $\varphi \in \boldsymbol{\theta}(\sigma'')$. Hence, $(\sigma, \sigma'') \in \boldsymbol{\rho}_i$.

To prove that for each $i \in N$ the relation $\boldsymbol{\rho}_i$ is connected, we must show that for all $\sigma, \sigma' \in \boldsymbol{V}$, either $(\sigma, \sigma') \in \boldsymbol{\rho}_i$ or $(\sigma', \sigma) \in \boldsymbol{\rho}_i$. This requires considerably harder and tedious work. If either $\boldsymbol{\theta}(\sigma)$ or $\boldsymbol{\theta}(\sigma')$ is empty, we are done immediately. Otherwise, the axioms *F3* and *F4* are heavily relied upon. For the remainder of the proof consider an arbitrary $i \in N$.

First we introduce the auxiliary notion of a *connecting path*, which we define as a sequence of vertices $\tau_0, \upsilon_0, \ldots, \tau_n, \upsilon_n$, or $\tau_0, \upsilon_0, \ldots \tau_{n-1}, \upsilon_{n-1}, \tau_n$ in $\boldsymbol{V}$ such that $\tau_0 = \epsilon$, and for each $0 \leqslant m \leqslant n$ both $\upsilon_m \in \boldsymbol{W}^{\tau_m}$ and $\tau_{m+1} \in T_{\upsilon_m}$. The latter requirement guarantees that, given the construction of $\mathfrak{T}_\Gamma$, each vertex τ_{m+1} was introduced to $\boldsymbol{V}$ in virtue of some formula $\neg[k]\chi$ being in $\boldsymbol{\theta}(\upsilon_m)$. Hence, we may assume that for each $0 < m \leqslant n+1$, there to be a $k \in N$ such that $\left\{\psi : \;\; [k]\psi \in \boldsymbol{\theta}(\upsilon_m)\right\} \subseteq \boldsymbol{\theta}(\tau_{m+1})$. Inspection of the various possible cases along with an easy inductive argument reveals that for each vertex σ in $\boldsymbol{W}$, there is a connecting path of which σ is the last element.

The argument now proceeds with a simultaneous induction on n and m, proving that for any two connecting paths $\sigma_0, \dots, \sigma_n$ and $\sigma'_0, \dots, \sigma'_m$ that either $(\sigma_n, \sigma'_m) \in \boldsymbol{\rho}_i$ or $(\sigma'_m, \sigma_n) \in \boldsymbol{\rho}_i$.

For $n = m = 0$, obviously, $\sigma_n = \sigma'_m = \epsilon$ and we are done immediately by reflexivity of $\boldsymbol{\rho}_i$. For the first inductive case consider two connecting paths $\sigma_0, \dots, \sigma_n, \sigma_{n+1}$ and $\sigma'_0, \dots, \sigma'_m$. If $n+1$ is odd, then $\sigma_{n+1} = \upsilon_{n/2}$ and $\sigma_n = \tau_{n/2}$. By definition of a connecting path then $\sigma_{n+1} \in \boldsymbol{W}^{\tau_{n/2}}$. If now, in this case, $\sigma_{n+1} = \sigma_n$, *i.e.*, if $\sigma_{n+1} \in \boldsymbol{T}$, then $\{\chi : \ [i]\chi \in \boldsymbol{\theta}(\sigma_n)\} \subseteq \boldsymbol{\theta}(\sigma_{n+1})$ as a consequence of axiom T_i and $\boldsymbol{\theta}(\sigma_{n+1})$ being maximal Λ-consistent. Otherwise in virtue of Fact A.4, $\{\chi : \ [\beta_{\sigma_{n+1}}]\chi \in \boldsymbol{\theta}(\sigma_n)\} \subseteq \boldsymbol{\theta}(\sigma_{n+1})$.

If, on the other hand, $n+1$ is even, then $\sigma_{n+1} = \tau_{n/2}$ and $\sigma_n = \upsilon_{(n/2)-1}$. In this case, there is a $k_{n+1} \in N$ such that $(\sigma_n, \sigma_{n+1}) \in \boldsymbol{\rho}_k$, *i.e.*, $\{\chi : \ [k_{n+1}]\chi \in \boldsymbol{\theta}(\sigma_n)\} \subseteq \boldsymbol{\theta}(\sigma_{n+1})$. Now let δ_{n+1} be the label in B defined as:

$$\delta_{n+1} \quad =_{df.} \quad \begin{cases} k_{n+1} & \text{if } n+1 \text{ is even,} \\ i & \text{if } n+1 \text{ is odd and } \sigma_{n+1} \in \boldsymbol{T}, \\ \beta_{\sigma_{n+1}} & \text{if } n+1 \text{ is odd and } \sigma_{n+1} \notin \boldsymbol{T}. \end{cases}$$

In general, we now have that $\{\chi : \ [\delta_{n+1}]\chi \in \boldsymbol{\theta}(\sigma_n)\} \subseteq \boldsymbol{\theta}(\sigma_{n+1})$.

By the induction hypothesis we may assume that $(\sigma_n, \sigma'_m) \in \boldsymbol{\rho}_i$ or $(\sigma'_m, \sigma_n) \in \boldsymbol{\rho}_i$. In the former case, assume that for some formula φ, $[i]\varphi \in \boldsymbol{\theta}(\sigma'_m)$ but $\varphi \notin \boldsymbol{\theta}(\sigma_{n+1})$ and consider an arbitrary formula ψ such that $[i]\psi \in \boldsymbol{\theta}(\sigma_{n+1})$. We show that $\psi \in \boldsymbol{\theta}(\sigma'_m)$. By maximal Λ-consistency of $\boldsymbol{\theta}(\sigma_{n+1})$, $[i]\psi \to \varphi \notin \boldsymbol{\theta}(\sigma_{n+1})$. This yields $[\delta_{n+1}]\big([i]\psi \to \varphi\big) \notin \boldsymbol{\theta}(\sigma_n)$. By maximal Λ-consistency and axiom $F4_{\delta_{n+1},i}$, then also that $[i]\big([i]\varphi \to \psi\big) \in \boldsymbol{\theta}(\sigma_n)$. Having assumed that $(\sigma_n, \sigma'_m) \in \boldsymbol{\rho}_i$, with the definition of $\boldsymbol{\rho}_i$, we have $[i]\varphi \to \psi \in \boldsymbol{\theta}(\sigma'_m)$. Subsequently, by maximal Λ-consistency and $[i]\varphi \in \boldsymbol{\theta}(\sigma_m)$, also $\psi \in \boldsymbol{\theta}(\sigma'_m)$.

If, on the other hand, $(\sigma'_m, \sigma_n) \in \boldsymbol{\rho}_i$, we prove again that $\psi \in \boldsymbol{\theta}(\sigma_m)$ for all formulas ψ such that $[i]\psi \in \boldsymbol{\theta}(\sigma_{n+1})$, provided that there be a formula φ such that $[i]\varphi \in \boldsymbol{\theta}(\sigma'_m)$ but $\varphi \notin \boldsymbol{\theta}(\sigma_{n+1})$. As before, $[i]\psi \to \varphi \notin \boldsymbol{\theta}(\sigma_{n+1})$ and consequently $[\delta_{n+1}]\big([i]\psi \to \varphi\big) \notin \boldsymbol{\theta}(\sigma_n)$. This time we had assumed that (σ'_m, σ_n). Hence, $[i][\delta_{n+1}]\big([i]\psi \to \varphi\big) \notin \boldsymbol{\theta}(\sigma'_m)$. By maximal Λ-consistency of $\boldsymbol{\theta}(\sigma'_m)$ and axiom $F3_{\delta_{n+1},i,i,i,i}$ then $[i][i]\big([i]\varphi \to \psi\big) \in \boldsymbol{\theta}(\sigma'_m)$. Two applications of axiom T_i, then give $[i]\varphi \to \psi \in \boldsymbol{\theta}(\sigma'_m)$. With $[i]\varphi$ having been assumed to be in $\boldsymbol{\theta}(\sigma'_m)$, the *desideratum* $\psi \in \boldsymbol{\theta}(\sigma'_m)$ follows. Figure 11 illustrates these arguments.

Since the argument for the second inductive case runs along analogous lines, we may conclude that each player i's the preference relation $\boldsymbol{\rho}_i$ is connected. □

A strategy profile of the extensive game $\boldsymbol{G}_\Gamma$ is then given by function mapping each internal vertex in $(\boldsymbol{V}_\Gamma, \boldsymbol{R}_\Gamma)$ onto a succeeding vertex. Not all strategy profiles of $\boldsymbol{G}_\Gamma$, however, are represented by a label in S. We will assume the label s in S to represent the strategy profile f_s that maps each internal vertex σ onto σs. *I.e.*, for each internal vertex σ in $\boldsymbol{V}_\Gamma$ and $s \in \boldsymbol{I}_\Gamma$ we have:

$$f_s(\sigma) \quad =_{df.} \quad \sigma s.$$

For all practical purposes we will identify f_s with s itself whenever no ambiguity is likely to arise.

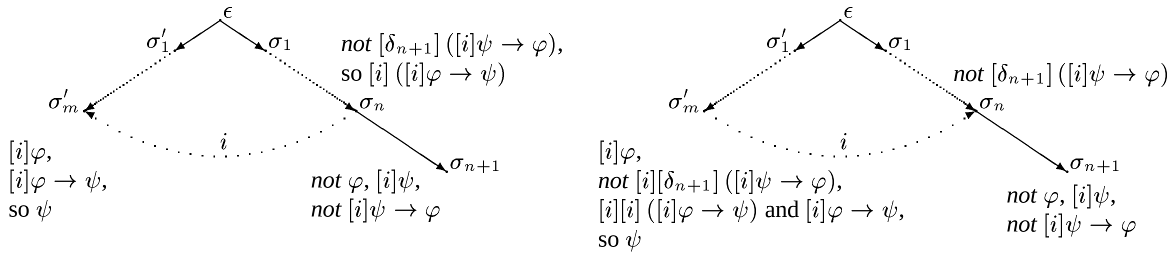

Figure 11.

Fact A.6. Let $x \in \boldsymbol{T}_\Gamma$. Then, for all $i \in N$ and $s \in S$:

$$\hat{s}_\emptyset(x) = \left\{xs^{\|N\|+1}\right\},$$
$$\hat{s}_i(x) = \left\{xs^{\|N\|+1}, xs^i s', xs^i l \in \boldsymbol{\Sigma}_\Gamma : \ s' \neq s \text{ and } l \in \omega\right\}.$$

Proof:
(*Sketch.*) On basis of the definition of $\hat{s}_N$ on page 289, and inspection of the construction of $\mathfrak{T}_\Gamma$ and $\boldsymbol{G}_\Gamma$, the following can be established:

$$\hat{s}_\emptyset(xs) = xs^{\|N\|+1}$$
$$\hat{s}_i(xs) = \left\{xs^{\|N\|+1}, xs^i s', xs^i l : \ s' \neq s \text{ and } l \in \omega\right\}.$$

Now recall that the additional player 0 had been assumed to be different from any of N. Since $\boldsymbol{P}(x) = 0$, for each $i \in N$ we have $\sigma \in s_i(x)$ if and only if $\sigma = s(x)$, *i.e.*, if $\sigma = xs$. By the definition of $\hat{s}_N$ on page 289 then $\hat{s}_i(x) = \hat{s}_i(xs)$ and $\hat{s}_\emptyset(x) = \hat{s}_\emptyset(xs)$, which give the desired result. □

We now define the model $\mathfrak{M}_\Gamma$ on the $\mathcal{L}(A, B)$-feasible game-frame $\mathfrak{F}_{\boldsymbol{G}_\Gamma}$.

Definition A.2. (Game-frame and game-model for $\boldsymbol{G}_\Gamma$)
Let Λ be an extensive game logic and Γ a Λ-consistent theory in a multi-modal matrix language $\mathcal{L}(A, B)$. The game-frame $\mathfrak{F}_{\boldsymbol{G}_\Gamma^\Lambda}$ on $\boldsymbol{G}_\Gamma^\Lambda$ is fixed by Definition 3.1, above, with $s \in S$ the label for the strategy profile f_s. Let $\mathfrak{F}_\Gamma^\Lambda$ denote $\mathfrak{F}_{\boldsymbol{G}_\Gamma^\Lambda}$. Let further the interpretation function $I_\Gamma \colon \boldsymbol{V}_\Gamma \to 2^A$, be defined by:

$$I_\Gamma(\sigma) \ =_{\textit{df.}} \ \left\{a \in A : \ a \in \boldsymbol{\theta}(\sigma)\right\}.$$

Then $\mathfrak{M}_\Gamma^\Lambda$ is the game-model $\left(\mathfrak{F}_\Gamma^\Lambda, I_\Gamma\right)$. When no confusion is likely to arise the superscript Λ is omitted.

The model $\mathfrak{M}_\Gamma$ can now be shown to satisfy the Λ-consistent theory Γ at the root node. In order to establish this, we prove something slightly stronger, *viz.*, that for each vertex v of $\mathfrak{M}_\Gamma$ the theory $\boldsymbol{\theta}_\Gamma(v)$, *i.e.*, the theory assigned to v by $\boldsymbol{\theta}_\Gamma$, contains exactly those formulas that are satisfied at v in $\mathfrak{M}_\Gamma$, provided that $\boldsymbol{\theta}_\Gamma(v)$ is non-empty.

Lemma A.1. (Truth lemma)
Let $\varGamma$ be a $\varLambda$-consistent theory in $\mathcal{L}(A, B)$. Consider both $\mathfrak{T}_\varGamma = (\boldsymbol{\varSigma}_\varGamma, \boldsymbol{\theta}_\varGamma)$ and $\mathfrak{M}_\varGamma$. Then for all vertices $\sigma \in \boldsymbol{W}_\varGamma$ and all formulas φ:

$$\mathfrak{M}_\varGamma, \sigma \Vdash \varphi \quad \text{iff} \quad \varphi \in \boldsymbol{\theta}_\varGamma(\sigma)$$

Proof:
Consider an arbitrary $x \in \boldsymbol{T}_\varGamma$ as well as an arbitrary $\sigma \in \boldsymbol{W}_\varGamma^x$. The proof is then by induction on φ.

For φ a propositional variable we are done immediately by the definition of $\mathfrak{M}_\varGamma$. Similarly, if φ is a Boolean combination, *i.e.*, $\varphi = \bot$, $\varphi = \neg\psi$ or $\varphi = \psi \wedge \chi$, maximal $\varLambda$-consistency of $\boldsymbol{\theta}(\sigma)$ takes care. Thus the modal cases remain.

Let $\varphi = [i]\psi$, for some $i \in N$. Assume φ to be the $n+1$-th formula in the enumeration, *i.e.*, $\varphi = \varphi_n$. First assume $[i]\psi \in \boldsymbol{\theta}(\sigma)$ and consider an arbitrary σ' such that $\sigma \boldsymbol{R}_i \sigma'$. By definition of $\boldsymbol{R}_i$ in $\mathfrak{M}_\varGamma$, also $\psi \in \boldsymbol{\theta}(\sigma')$. Consequently $\boldsymbol{\theta}(\sigma')$ is not empty and so $\sigma' \in \boldsymbol{W}_\varGamma$. Therefore, the induction hypothesis is applicable, rendering $\mathfrak{M}_\varGamma, \sigma' \Vdash \psi$. Having chosen σ' arbitrarily, $\mathfrak{M}_\varGamma, \sigma \Vdash [i]\psi$ follows. For the opposite direction, assume $[i]\psi \notin \boldsymbol{\theta}(\sigma)$. With maximal $\varLambda$-consistency, then $\neg[i]\psi \in \boldsymbol{\theta}(\sigma)$. Hence, for t_n^σ, we have $\{\neg\psi\} \cup \{\chi : [i]\chi \in \boldsymbol{\theta}(\sigma)\} \subseteq \boldsymbol{\theta}(t_n^\sigma)$. By definition of $\boldsymbol{R}_i$ in $\mathfrak{M}_\varGamma$, then $\sigma \boldsymbol{R}_i t_n^\sigma$. Again, $t_n^\sigma \in \boldsymbol{W}_\varGamma$ and with the induction hypothesis $\mathfrak{M}_\varGamma, t_n^\sigma \nVdash \psi$. Accordingly, $\mathfrak{M}_\varGamma, \sigma \nVdash [i]\psi$.

Let $\varphi = [\hat{s}_i]\psi$, for some $i \in N$ and $s \in S$. Again we distinguish the cases in which σ is a leaf from the one in which σ and x are identical. First, assume the former. Then, σ itself is the only element of $\hat{s}_i(\sigma)$, *i.e.*, $\hat{s}_i(\sigma) = \{\sigma\}$. Observe further that in virtue of axiom $F2_{\beta_\sigma, \hat{s}_i}$, $[\beta_\sigma]\big([\hat{s}_i]\psi \leftrightarrow \psi\big)$ is an element of $\boldsymbol{\theta}(x)$. With Fact A.4 then $[\hat{s}_i]\psi \leftrightarrow \psi \in \boldsymbol{\theta}(\sigma)$. Now consider the following equivalences:

$$\begin{aligned}
[\hat{s}_i]\psi \in \boldsymbol{\theta}(\sigma) \quad & \text{iff}_{[\hat{s}_i]\psi \leftrightarrow \psi \in \boldsymbol{\theta}(\sigma)} && \psi \in \boldsymbol{\theta}(\sigma) \\
& \text{iff}_{\textit{i.h.}} && \mathfrak{M}_\varGamma, \sigma \Vdash \psi \\
& \text{iff}_{\hat{s}_i(\sigma) = \{\sigma\}} && \text{for all } \sigma' \text{ such that } \sigma' \in \hat{s}_i(\sigma) : \mathfrak{M}_\varGamma, \sigma' \Vdash \psi \\
& \text{iff} && \text{for all } \sigma' \text{ such that } \sigma \boldsymbol{R}_{\hat{s}_i(\sigma)} \sigma' : \mathfrak{M}_\varGamma, \sigma' \Vdash \psi \\
& \text{iff} && \mathfrak{M}_\varGamma, \sigma \Vdash [\hat{s}_i]\psi.
\end{aligned}$$

If, on the other hand, σ and x are identical, we reason as follows. First assume that $[\hat{s}_i] \in \boldsymbol{\theta}(x)$ and consider an arbitrary σ' such that $\sigma \boldsymbol{R}_{\hat{s}_i} \sigma'$. It suffices to prove that $\mathfrak{M}_\varGamma, \sigma' \Vdash \psi$. By axiom $F1_{\hat{s}_i, \hat{s}_\emptyset}$, also $[\hat{s}_\emptyset]\psi \in \boldsymbol{\theta}(x)$. Inspection of Fact A.6 reveals that $[\beta_{\sigma'}]\psi \in \boldsymbol{\theta}(x)$ and hence, by Fact A.4 we have $\psi \in \boldsymbol{\theta}(\sigma')$. By the induction hypothesis follows that $\mathfrak{M}_\varGamma, \sigma \Vdash \psi$ and with σ' having been chosen arbitrarily, $\mathfrak{M}, x \Vdash [\hat{s}_i]\psi$.

For the opposite direction, assume that $[\hat{s}_i]\psi \notin \boldsymbol{\theta}(x)$. Then, by maximal $\varLambda$-consistency, $\neg[\hat{s}_i]\psi \in \boldsymbol{\theta}(x)$. Without loss of generality we may assume $\neg[\hat{s}_i]\psi$ to be the $n+1$-st element of the enumeration of the formulas we considered. Now let $\sigma^* =_{\textit{df.}} x s^i n$. Then, $\sigma^* \in \hat{s}_i(x)$ and *a fortiori* also $x \boldsymbol{R}_{\hat{s}_i} \sigma^*$. Moreover, by construction $\neg\psi \in \boldsymbol{\theta}(\sigma^*)$. By the induction hypothesis, then, $\mathfrak{M}_\varGamma, \sigma^* \Vdash \neg\psi$ and, consequently, $\mathfrak{M}, \sigma^* \nVdash \psi$, which suffices for a proof.

Let $\varphi = [\hat{s}_\emptyset]$, for some $s \in S$. Let $\sigma^{**} = x s^{\|N\|+1}$. Again we distinguish between σ being a leaf in $\boldsymbol{L}^x$ and σ being x itself, dealing with the latter case first.

First assume that $[\hat{s}_\emptyset]\psi \in \boldsymbol{\theta}(x)$. By Fact A.4 then $\psi \in \boldsymbol{\theta}(\sigma^{**})$. The induction hypothesis subsequently yields $\mathfrak{M}_\varGamma, \sigma^{**} \Vdash \psi$. From Fact A.6 we learn that σ^{**} is the only element in $\hat{s}_\emptyset(x)$ and, hence, we may conclude that $\mathfrak{M}_\varGamma, x \Vdash [\hat{s}_\emptyset]\psi$.

For the opposite direction assume $[\hat{s}_\emptyset]\psi \notin \boldsymbol{\theta}(x)$. By maximal Λ-consistency of $\boldsymbol{\theta}(x)$ then both $\neg[\hat{s}_\emptyset]\psi \in \boldsymbol{\theta}(x)$ and $\langle\hat{s}_\emptyset\rangle\neg\psi \in \boldsymbol{\theta}(x)$. In virtue of axiom $D!_{\hat{s}_\emptyset}$, then $[\hat{s}_\emptyset]\neg\psi \in \boldsymbol{\theta}(x)$. By Fact A.4 it can readily be established that $\neg\psi \in \boldsymbol{\theta}(\sigma^{**})$ and with maximal Λ-consistency of $\boldsymbol{\theta}(\sigma^{**})$ also $\psi \notin \boldsymbol{\theta}(\sigma^{**})$. The induction hypothesis is applicable and so $\mathfrak{M}_\Gamma, \sigma^{**} \nVdash \psi$. Fact A.6 guarantees that $\sigma^{**} \in \hat{s}_\emptyset(x)$ and *a fortiori* $x\boldsymbol{R}_{\hat{s}_\emptyset}\sigma^{**}$. Hence, eventually, $\mathfrak{M}_\Gamma, x \nVdash [\hat{s}_\emptyset]\psi$, which we had to prove.

Now suppose σ to be a leaf in $\boldsymbol{L}^x$. Then $\hat{s}_\emptyset(\sigma) = \{\sigma\}$. Observe that in virtue of axiom $F5_{\beta_\sigma,\hat{s}_\emptyset}$ we have that $[\beta_\sigma]\big([\hat{s}_\emptyset]\varphi \leftrightarrow \varphi\big) \in \boldsymbol{\theta}(x)$. By Fact A.4 then also $[\hat{s}_\emptyset]\varphi \leftrightarrow \varphi \in \boldsymbol{\theta}(\sigma)$. Now consider the following equivalences:

$\mathfrak{M}_\Gamma, \sigma \Vdash [\hat{s}_\emptyset]\psi$	iff	for all $\sigma' \in \boldsymbol{V}$: $\sigma\boldsymbol{R}_{\hat{s}_\emptyset}\sigma'$ implies $\mathfrak{M}_\Gamma, \sigma' \Vdash \psi$
	iff	for all $\sigma' \in \hat{s}_\emptyset(\sigma) : \mathfrak{M}_\Gamma, \sigma' \Vdash \psi$
	$\text{iff}_{\hat{s}_\emptyset(\sigma)=\{\sigma\}}$	$\mathfrak{M}_\Gamma, \sigma \Vdash \psi$
	$\text{iff}_{\textit{i.h.}}$	$\psi \in \boldsymbol{\theta}(\sigma)$
	$\text{iff}_{[\hat{s}_\emptyset]\varphi \leftrightarrow \varphi \in \boldsymbol{\theta}(\sigma)}$	$[\hat{s}_\emptyset] \in \boldsymbol{\theta}(\sigma)$.

This concludes the proof. □

Completeness of the minimal extensive game logic F with respect to the semantics in terms of game-frames now follows as a corollary of this last result.

Theorem A.1. (Completeness of F)
Let Γ be a theory and φ a formula in a multi-modal matrix language $\mathcal{L}(A, B)$. Then:

$$\Gamma \vdash_{\mathsf{F}} \varphi \quad \text{iff} \quad \Gamma \models_{\mathsf{F}} \varphi.$$

Proof:
The left-to-right direction is taken care of by Proposition 4.1, above. For the right-to-left direction it suffices to prove that there is a model on a game-frame for each F-consistent theory Γ in $\mathcal{L}(A, B)$. The construction of $\mathfrak{M}^{\mathsf{F}}_\Gamma$ for F, as defined above, provides such a model. For, proving the routine contrapositive, $\Gamma \nvdash_{\mathsf{F}} \varphi$ implies $\Gamma \cup \{\neg\varphi\} \nvdash_{\mathsf{F}} \bot$. Then $\mathfrak{M}^{\mathsf{F}}_{\Gamma\cup\{\neg\varphi\}}$ exists and, by Lemma A.1, $\mathfrak{M}^{\mathsf{F}}_{\Gamma\cup\{\neg\varphi\}}, \epsilon \Vdash \Gamma \cup \{\neg\varphi\}$ F. Then, also $\mathfrak{M}^{\mathsf{F}}_{\Gamma\cup\{\neg\varphi\}}, \epsilon \Vdash \Gamma$ and $\mathfrak{M}^{\mathsf{F}}_{\Gamma\cup\{\neg\varphi\}}, \epsilon \nVdash \varphi$, yielding $\Gamma \nvDash_{\mathsf{F}} \varphi$. □

Completeness for the extensive game logics $\mathsf{F5}_{s,i}$ and $\mathsf{F5}^N_s$ (for particular labels s and i and N the whole set of player-labels) can be obtained in a similar fashion. The validity of the argument, however, rests on the fact that the models $\mathfrak{M}^{\mathsf{F5}_{s,i}}_\Gamma$ and $\mathfrak{M}^{\mathsf{F5}^N_i}_\Gamma$ belong to the appropriate classes of game-models. *I.e.*, for each $\mathsf{F5}_{i,s}$-consistent theory Γ the model $\mathfrak{M}^{\mathsf{F5}_{i,s}}_\Gamma$ is based on an extensive game in which the strategy profile s contains a best response for player i. Similarly, in the extensive game underlying a model $\mathfrak{M}^{\mathsf{F5}^N_s}_\Gamma$ for a $\mathsf{F5}^N_s$-consistent theory Γ, there is a label in N for each interested player and the strategy profile denoted by the label s is a subgame perfect Nash equilibrium. Hence, we have the following results.

Theorem A.2. The logic $\mathsf{F5}_{i,s}$ is sound and complete with respect to the class of game-models built on game-frames in which s is a subgame perfect best response for player i.

Proof:
Soundness follows from Theorem 3.1, on page 293 above. For completeness the proof is as in Theorem A.1, be it that it should also be shown that for any $\mathsf{F5}_{i,s}$-consistent theory, s is a subgame perfect best response for player i in the extensive game $\boldsymbol{G}_\Gamma$ that is defined in the course of the construction of the model $\mathfrak{M}_\Gamma^{\mathsf{F5}_{i,s}}$. In virtue of Proposition 3.1 on page 293, we may restrict ourselves to showing that $\mathfrak{F}_\Gamma$ is $(\hat{s}_i, i, \hat{s}_\emptyset)$-Euclidean. So consider an arbitrary $\mathsf{F5}_{i,s}$-consistent theory Γ and equally arbitrary vertices $\sigma, \sigma', \sigma'' \in \boldsymbol{V}_\Gamma$ such that $\sigma' \in \hat{s}_i(\sigma)$ and $\sigma'' \in \hat{s}_\emptyset(\sigma)$. We prove that $(\sigma', \sigma'') \in \boldsymbol{\rho}_i$. So consider an arbitrary formula φ and assume that $[i]\varphi \in \boldsymbol{\theta}(\sigma')$. It suffices to demonstrate that $\varphi \in \boldsymbol{\theta}(\sigma'')$. In case σ is a leaf node, then σ, σ' and σ'' are all identical, and we are done immediately by T_i, *i.e.*, by reflexivity of $\boldsymbol{\rho}_i$ in $\boldsymbol{G}_\Gamma$. So, for the remainder of the proof, we may assume σ to be an internal vertex. Without loss of generality, we may assume $\sigma = xs^m$ and $\sigma'' = xs^{\|N\|+1}$, where $x \in \boldsymbol{T}$ and $0 \leqslant m \leqslant \|N\|$. If it so happens that $i < m$, then again σ' and σ'' are identical — as can easily be established — and we are done by axiom T_i. Otherwise, observe that, having assumed $[i]\varphi \in \boldsymbol{\theta}(\sigma')$, it holds that $\langle\beta_{\sigma'}\rangle[i]\varphi \in \boldsymbol{\theta}(x)$. Hence, in virtue of $\langle\hat{s}_i\rangle[i]\varphi \to [\hat{s}_\emptyset]\varphi$ being an axiom of $\mathsf{F5}_{i,s}$ and maximal $\mathsf{F5}_{i,s}$-consistency of $\boldsymbol{\theta}(x)$, also $[\hat{s}_\emptyset]\varphi \in \boldsymbol{\theta}(x)$. Eventually, by definition of $\boldsymbol{\theta}\left(xs^{\|N\|+1}\right)$ also $\varphi \in \boldsymbol{\theta}\left(xs^{\|N\|+1}\right)$. This concludes the proof. $\square$

Theorem A.3. The logic $\mathsf{F5}_s^N$ in $\mathcal{L}(A, B)$ is sound and complete with respect to the class of game-frames built on games in which s is a subgame perfect Nash equilibrium and in which there is a label in N for each interested player.

Proof:
Soundness is again immediate by Theorem 3.1. For completeness, it suffices to show that, for each maximal $\mathsf{F5}_s^N$-consistent theory, s is a subgame perfect Nash equilibrium in the extensive game $\boldsymbol{G}_\Gamma$, which is built in the construction of the model $\mathfrak{M}_\Gamma^{\mathsf{F5}_s^N}$. Observe that for each label $i \in N$, the axiom $5_{\hat{s}_i, i, \hat{s}_\emptyset}$ is derivable in $\mathsf{F5}_s^N$ in $\mathcal{L}(A, B)$. Hence, in a similar manner as in Theorem A.2, we can show that for each *interested* player the strategy profile s is a best response for player i in $\boldsymbol{G}_\Gamma$. This leaves the mystery player 0. His, her or its preference relation is universal by construction. As a consequence, each strategy profile is a best response for 0. As this is the case in particular for s, we are done. $\square$

The extensive games $\boldsymbol{G}_\Gamma^\Lambda$ for Λ-consistent theories have some noteworthy properties in common. In particular, the depth of the game $\boldsymbol{G}_\Gamma^\Lambda$ — *i.e.*, the length of the longest path in the game-tree connecting the root to a leaf — does not exceed the number of player labels in $\mathcal{L}(A, B)$ plus two, *i.e.*, $\|N\| + 2$. Moreover, the players are assumed to play in a fixed order, and on each path in the game-tree from the root to a leaf, each player represented by a label in N moves at most once and any other player at most twice. Also, the number of players in each game $\boldsymbol{G}_\Gamma$ is always one greater than the number of labels in N.

Corollary A.1. Let G be an extensive game of perfect information G with N the set of players and let Γ be a theory in a feasible multi-modal matrix language $\mathcal{L}(A, B)$ with N the set of player labels. Assume further that Γ is satisfiable in a model on $\mathfrak{F}_G$. Then there is an extensive game of perfect information G' a game-model on which also satisfies Γ and which game tree has a maximal depth of $\|N\| + 2$ and which players number $\|N\| + 1$. Moreover, in each play of G' each player represented by a label in N moves at most once and any other player at most twice.

Proof:
(*Sketch.*) Since Γ is satisfiable in a model on $\mathfrak{F}_G$, by Theorem A.1, F-consistency of Γ follows. Construct the extensive game $\boldsymbol{G}_\Gamma^{\mathsf{F}}$. The truth lemma A.1 assures that Γ is satisfied at the root node of $\boldsymbol{G}_\Gamma^{\mathsf{F}}$. Moreover, $\boldsymbol{G}_\Gamma^{\mathsf{F}}$ can be seen to comply to the requirements as stated in the corollary. □

This corollary says, roughly, that one can confine one's attention to games of a limited depth when studying finite extensive games with respect to properties of extensive games expressible in the respective multi-modal matrix language.

References

[1] Aumann, R. J.: Game Theory, in: *Game Theory* (J. Eatwell, M. Milgate, P. Newman, Eds.), The New Palgrave, Macmillan, London and Basingstoke, 1997, 1–54.

[2] Baltag, A.: A Logic for Games, *Proceedings of the ILLC Workshop on Logic and Games, Held in Amsterdam, November 19-20, 1999* (M. Pauly, A. Baltag, Eds.), number PP-1999-25 in ILLC Prepublications Series, ILLC, Amsterdam, November 1999 1999.

[3] Benthem, J.F.A.K. van: *Logic and Games. Lecture Notes,* ILLC, Amsterdam, 2001.

[4] Blackburn, P., de Rijke, M., Venema, Y.: *Modal Logic*, vol. 53 of *Cambridge Tracts in Theoretical Computer Science,* Cambridge U.P., Cambridge, 2001.

[5] Bonanno, G.: *Branching Time Logic, Perfect Information Games and Backward Induction*, Department of Economics, University of California, 1998.

[6] Boutilier, C., Shoham, Y., Wellman, M. P.: Economic Principles of Multi-Agent Systems, *Artificial Intelligence*, **94**, 1997, 1–6.

[7] Fudenberg, D., Tirole, J.: *Game Theory,* The MIT Press, Cambridge, Mass., 1991.

[8] Goldblatt, R.: *Logics of Time and Computation*, vol. 7 of *CSLI Lecture Notes*, 2nd edition, CSLI Publications, Stanford, 1992.

[9] Harel, D., Kozen, D., Tiuryn, J.: *Dynamic Logic,* MIT Press, Cambridge, Mass. & London, 2001.

[10] Harrenstein, B. P., van der Hoek, W., Meyer, J.-J., Witteveen, C.: Subgame Perfect Nash-Equilibria in Dynamic Logic, *Proceedings of the ILLC Workshop on Logic and Games, Held in Amsterdam, November 19-20, 1999* (M. Pauly, A. Baltag, Eds.), number PP-1999-25 in ILLC Prepublications Series, ILLC, Amsterdam, 1999.

[11] Harrenstein, B. P., van der Hoek, W., Meyer, J.-J., Witteveen, C.: A Modal Interpretation of Nash-Equilibria and Some Related Concepts, *Proceedings of the Fourth Conference on Logic and the Foundations of Game and Decision Theory (LOFT4), June 30–July 2, 2000* (B. Giacomo, E. Colombatto, W. van der Hoek, Eds.), ICER, Torino, 2000.

[12] Harrenstein, B. P., van der Hoek, W., Meyer, J.-J., Witteveen, C.: On Modal Logic Interpretations of Games, *ECAI 2002, 15th European Conference on Artificial Intelligence, July 21–26 2002, Lyon, France* (F. van Harmelen, Ed.), 77, IOS Press, Amsterdam, 2002, ISSN 0922-6389.

[13] Jongh, D. de, Veltman, F.: *Intensional Logics,* FdW, University of Amsterdam, Amsterdam, 1995.

[14] Kozen, D.: Results in the Propositional μ-Calculus, *Theoretical Computer Science*, **27**, 1983, 333–354.

[15] Kreps, D. M.: Nash Equilibrium, in: *Game Theory* (J. Eatwell, M. Milgate, P. Newman, Eds.), The New Palgrave, Macmillan, London and Basingstoke, 1997, 157–177.

[16] Kuhn, H. W.: Extensive Games and the Problem of Information, in: *Contributions to the Theory of Games, Volume II* (H. W. Kuhn, A. W. Tucker, Eds.), number 28 in Annals of Mathematics Studies, Princeton U.P., Princeton, N.J., 1953, 1–54, Reprinted in [17].

[17] Kuhn, H. W., Ed.: *Classics in Game Theory*, Princeton U.P., Princeton, N.J., 1997.

[18] Luce, R. D., Raiffa, H.: *Games and Decisions. Introduction and Critical Survey*, John Wiley & Sons, New York, 1957.

[19] Myerson, R. B.: *Game Theory, Analysis of Conflict*, Harvard U.P., Cambridge, Mass., 1991.

[20] Nash, J.: Equilibrium Points in n-Person Games, *PNAS*, **36**, 1950, 48–49, Reprinted in [17].

[21] Nash, J.: Non-Cooperative Games, *Annals of Mathematics*, **54**, 1951, 286–295, Reprinted in [17].

[22] Neumann, J. von, Morgenstern, O.: *Theory of Games and Economic Behavior*, Princeton U.P., Princeton, N.J., 1944.

[23] Osborne, M. J., Rubinstein, A.: *A Course in Game Theory*, MIT Press, Cambridge, Mass., 1994.

[24] Parikh, R.: Logics of Knowledge, Games and Dynamic Logic, in: *Foundations of Software Technology and Theoretical Computer Science* (M. Joseph, R. Shyamasundar, Eds.), number 181 in Lecture Notes in Computer Science, Springer, New York, 1984, 202–222.

[25] Parikh, R.: The Logic of Games and its Applications, in: *Topics in the Theory of Computation, FCT'83, Borgholm, Sweden, August 21-27, 1983* (M. Karpinski, J. van Leeuwen, Eds.), number 14 in Annals of Discrete Mathematics, North-Holland, Amsterdam, 1985, 111–140.

[26] Pauly, M.: *Logic for Social Software*, Ph.D. Thesis, Institute for Logic, Language and Information, Amsterdam, 2001.

[27] Popkorn, S.: *First Steps in Modal Logic*, Cambridge U.P., Cambridge, 1994.

[28] Schelling, T. C.: *The Strategy of Conflict*, Harvard U.P., Cambridge, Mass., 1960.

[29] Selten, R.: Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit, *Zeitschrift für die gesamte Staatswissenschaft*, **121**, 1965.

[30] Stirling, C.: Games and Modal Mu-Calculus, in: *Proceedings of Second International Workshop on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'96)* (T. Margaria, B. Steffen, Eds.), number 1055 in Lecture Notes in Computer Science, Springer, New York, 1996, 298–312.

[31] Zermelo, E.: Über eine Anwendung der Mengenlehre auf die Theorie des Schachspiels, *Proceedings of the Fifth International Congress of Mathematicians* (E. W. Hobson, A. E. H. Love, Eds.), II, Cambridge U.P., Cambridge, 1913.