On Modal Logic Interpretations of Games

Paul Harrenstein[1] and **Wiebe van der Hoek**[1,2] and **John-Jules Meyer**[1] and **Cees Witteveen**[3]

Abstract. Multi-agent environments comprise decision makers whose deliberations involve reasoning about the expected behavior of other agents. Apposite concepts of rational choice have been studied and formalized in game theory and our particular interest is with their integration in a logical specification language for multi-agent systems. This paper concerns the logical analysis of the game-theoretical notions of a (subgame perfect) Nash equilibrium and that of a (subgame perfect) best response strategy. Extensive forms of games are conceived of as Kripke frames and a version of Propositional Dynamic Logic is employed to describe them. We show how formula schemes of our language characterize those classes of frames in which the strategic choices of the agents can be said to be Nash-optimal. Our analysis focuses on extensive games of perfect information without repetition.

1 Introduction

Agents can be thought of as systems that are capable of reasoning about their own and other agents' knowledge, preferences, goals, future and past actions. A successful framework to formally reason about such agent systems is provided by *modal logic* (see ([11] or ([7])) of which the operators model various mental attitudes and also the dynamics of such a system. Although the approaches mentioned enable one to reason about the adoption and persistence of goals and intentions, they do not explain where such intentions come from. As an agent may be confronted with several, mutually exclusive, ways how to act, however, decision making and intention formation is imperative. Which action an agent eventually performs may very well depend on his beliefs concerning the other agents' actions and their responses to his actions. Since *game theory* is devoted to the study of such reasoning mechanisms and the associated notion of strategic rationality, many of its concepts are relevant to the study of multi-agent systems.

The emphasis of this paper is on the incorporation of some important game-theoretical notions in Propositional Dynamic Logic (cf., [5]). Although our work is still of a purely theoretical nature, we believe it to be conductive to the development of a comprehensive logical framework in which multi-agent systems can be described, specified and reasoned about. In that sense, our work chimes in well with a direction set out by the communities of logic and game theory to integrate the two research paradigms, thereby providing logical tools for reasoning about rationality and decision making on the one hand, but also importing game-theoretic notions into the realm of logic, on the other. Our work is best seen as a proponent of the first, with the overall goal to incorporate game-theoretic arguments in logical formalisms that model BDI-like mental attitudes.

The key idea behind our approach is the similarity between the semantics of modal logic (Kripke models), on the one hand, and games in extensive form, on the other. We propose to judge such a game tree as a structure for dynamic logic, in which agents can perform programs (moves) to achieve their goals (maximal utility). Along with the closely related concept of a best response strategy our investigations focus on the solution concept of a Nash equilibrium. We also deal with their subgame perfect varieties. We show how the 'modal counterparts' of such concepts look like. As such, we pave the way for a more thorough analyses which enables to study these game-theoretic concepts within a BDI- or KARO-like framework. Again, such a unification harvests in two ways: the logical approaches to agents are enriched with a game-theoretic perspective, but also, well known constructs from dynamic logic, for example, may give rise to new concepts (like operations on games) in game theory (cf., [2]).

2 Game Theoretical Notions

The investigations of this paper concern finite games in extensive form with perfect information. A (pure) *strategy* for a game consists of a complete plan for a player i how to play that game. *Strategy profiles*, denoted by σ, combine strategies, one for each player. A strategy profile determines for each node a unique outcome, though not necessarily for each node the same one.

Example 1 *Consider the two-person game in extensive form as depicted in Figure 1. Let RL denote the strategy for player i_1 that consists in his going right at v_0 and going left at v_3. RL can be conceived of as the function that maps v_0 onto v_2 and v_3 onto z_1. The*

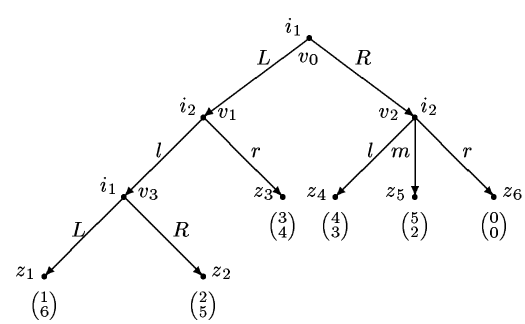

Figure 1. Example of a game in extensive form

[1] Institute of Information & Computing Science, Utrecht University
[2] Computer Science Department, The University of Liverpool
[3] Faculty of Information Technology and Systems, Delft University of Technology

pair of strategies (RL, ll) — where ll is the strategy for player i_2 which prescribes her to go left at both v_1 and v_2 — denotes a strategy profile and determines the outcome z_4, and granting payoffs of 4 *and* 3 *to i_1 and i_2, respectively.*

Whether a strategy is a best response for a player is relative to the strategies the other players adopt, i.e., to a strategy profile. Assuming that play commences at the root node, a strategy profile σ is said to contain a *best response for player i_k*, if i_k cannot increase her payoff by playing another strategy available to her when the other players stick to their strategies as specified in σ. A strategy profile is a *Nash-equilibrium* if none of the players can increase her payoff by unilaterally playing another strategy. Equivalently, a Nash equilibrium could be characterized as a strategy profile which contains a best response strategy for each player (cf. [8], p. 98).

It has been argued that Nash equilibria do not in general do justice to the sequential structure of an extensive game. In our example, (LR, lr) is, along with (RL, ll), (RL, rl), (RR, ll) and (RR, rl), a Nash equilibrium. This is, however, dependent on the fact that i_2 going right at v_2 minimizes i_1's pay-off rather than that it maximizes that of i_2. Player i_2, as it were, threatens to go right at v_2 if i_1 goes right at v_0. Player i_1, however, need not take this threat seriously if the sequentiality of the game is taken into account. The node v_2 will be reached only if i_1 moved right at v_0 at a previous state of the game. Once in v_2, strategic rationality prescribes i_2 to move to z_4 rather then go right to z_6. As there is nothing in the description of the game committing i_2 to move to z_6 in v_2, the strategy profile (LR, lr) should be ruled out as a rational alternative. This is a manifestation of the phenomenon that a strategy profile contains instructions for the players how to act in nodes that they preclude ever to be reached in the course of the game and in some cases allows for "irrational" moves off the equilibrium path.

A refinement of the solution concept of Nash equilibrium that meets this objection is achieved by requiring Nash equilibria to be *subgame perfect*. In extensive form, a subgame can be conceived of as a cutting of the game tree, which results in another game in extensive form. A strategy is a *subgame perfect* best response strategy for a player relative to some strategy profile σ if it is a best response strategy with respect to σ in all its subgames.

A strategy profile determines a unique outcome. By deviating unilaterally, a player can force several outcomes to come about by choosing her strategy. The one guaranteeing her the highest outcome is her best response strategy with respect to the respective strategy profile. These outcomes can be represented graphically by the leaf nodes of the game tree from which are removed all edges that do not comply with the strategies of the other players as laid down in the strategy profile. Such a reduced tree we call a player's *strategy search space* with respect to a strategy profile. In our example the strategy search space for i_1, given a strategy profile containing i_2's strategy ll, is obtained from Figure 1 by removing the edges to z_3, z_5 and z_6.

Game trees, being graphs, correspond to Kripke structures and as such they can be described by means of the language of Propositional Dynamic Logic (PDL). The nodes of the game tree represent the states of the frame and the edges define the accessibility relation. A strategy for a player i is identified with the graph of a function from the nodes at which i can move to successor nodes. A strategy profile combines strategies of the individual players and as such it is the graph of a function on the internal nodes of the game tree. In this manner, strategies, strategy profiles and strategy search spaces can be represented by programs of dynamic logic.

Fundamental to the present analysis is that frames in which the program representing a strategy profile σ contains a (subgame perfect) best response strategy for some player or is a (subgame perfect) Nash-equilibrium, possess certain structural properties which are expressible in PDL. The objective of this paper is to specify formally which constraints a frame satisfies if the strategy program corresponds to a strategy profile that is a (subgame perfect) Nash-equilibrium or incorporates a (subgame perfect) best response strategy. We show how formula schemes of PDL characterize the frames satisfying these structural properties.

So far the concepts of game theory relevant to this paper have only been presented in a rather informal fashion. Now, we give a formal account in which we go a long way in following Bonanno's (cf. [4]). A game in extensive form with perfect information without repetition is identified with a a tuple $(V, \prec, Z, N, \iota, u)$, where V is a finite set of vertices and $\prec$ a relation on V, representing the possible moves at each vertex. The pair $(V, \prec)$ is a non-trivial, irreflexive, finite tree. Furthermore, Z is the set of leaves of the tree, and N is a finite set of players. The function ι assigns a player to each internal node of the game tree and is supposed to be surjective (onto). Finally, we have the *utility function* $u : N \times Z \to \mathbb{R}$ assign a pay-off to each player in N in each vertex, rather than just at the leaves as is costumary. We write $z \leq_i z'$ for $u(i)(z) \leq u(i)(z')$.

As a notational convention, we use v_0 to denote the root and define for each i, $V_i := \{v \in V : \ \iota(v) = i\}$.

Definition 2 *(Strategies and Strategy Profiles) Let $G = (V, \prec, Z, N, \iota, u)$ be an extensive game. A* strategy profile *for G is a total function $\sigma : V \backslash Z \to V$ such that for all $v \in V$, $v \prec \sigma(v)$. The set of strategy profiles we denote by Σ. For $W \subseteq V$, we have $\sigma \sim_W \sigma'$ iff for all $w \in W : \sigma(w) = \sigma'(w)$, writing $\sigma \sim_{-i} \sigma'$ for $\sigma \sim_{V \backslash V_i} \sigma'$.*

A strategy *for a player $i \in N$ is the restriction of a strategy profile σ to V_i, i.e. , $\sigma |_{V_i}$.*

Each strategy profile σ determines a unique outcome in the sense that if all players stick to the strategies in σ, the game terminates in precisely one final stage. As such, a strategy profile σ gives rise to a function that maps each internal node v to the leaf σ fixes as its outcome when playing the game is started in v. To capture this notion we define for each strategy profile σ, the function $\hat{\sigma} : V \to Z$ by:

$$\hat{\sigma}(v) \ = \ \begin{cases} v & \text{if } v \text{ is a leaf} \\ \hat{\sigma}(\sigma(v)) & \text{otherwise.} \end{cases}$$

We now give formal definitions of the game theoretical notions of a best response strategy relative to a strategy profile and a Nash-equilibrium, as well as their subgame perfect, or s-p., variations:

Definition 3 *Let $G = (V, \prec, Z, N, \iota, u)$ be a game on an extensive form and $\sigma \in \Sigma_G$. Let further v_0 be the root of G and let i range over N, v over V and σ over Σ. Then define:*

(1) *σ comprises a* best response strategy for *i iff*
for all $\sigma' \in \Sigma$: $\sigma \sim_{-i} \sigma'$ implies $\hat{\sigma}'(v_0) \leq_i \hat{\sigma}(v_0)$

(2) *σ is a* Nash-equilibrium *iff*
for all $i \in N$, σ comprises a best response for i

(3) *σ comprises an* s.p.-best response strategy for *i iff*
for all $v \in V, \sigma' \in \Sigma$: $\sigma \sim_{-i} \sigma'$ implies $\hat{\sigma}'(v) \leq_i \hat{\sigma}(v)$

(4) *σ comprises an* s.p.-Nash-equilibrium *iff*
for all $i \in N$, σ is a s.p. best response for i.

3 Logic: Syntax & Semantics

Being graphs, game trees can straightforwardly be correlated with Kripke structures, for which description modal languages are available. The analyses of this paper are conducted in languages for PDL augmented by a separate set of labelled modal operators $\{[i]\}_{i\in\Lambda}$. Reinforced thus, we argue the language gains expressive power with respect to the players' preference orderings on the possible outcomes as they are determined by the payoff structure of the corresponding game.

The resulting logical language is a multi-modal dynamic language containing the usual Boolean operations "$\neg$" (*negation*) and "$\wedge$" (conjunction) and as program connectives, ";", "$\cup$" and "$*$" (iteration) as well as a program forming operation on formulae "?" (test).

Definition 4 *(Syntax) Let Φ_0 be a countable set of propositional variables (typical element A), Π_0 be a set of atomic programs (typical element a) and Λ a set of labels (typical element i). The set of formulae, Φ (typical element φ) and the set of programs Π (typical element α) are generated by the following grammar:*

$$\begin{array}{lcl} \varphi & ::= & A \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid [\pi]\varphi \mid [i]\,\varphi \\ \pi & ::= & a \mid \pi_1;\pi_2 \mid \pi_1 \cup \pi_2 \mid \pi^* \mid \varphi? \end{array}$$

We will assume for each extensive form $G = (V, \prec, Z, N, \iota, u)$ a language L_G with Λ at least as large as N and Π_0 at least as large as $N \cup \Sigma$ i.e., $|N| \le |\Lambda|$ and $|N \cup \Sigma| \le |\Pi_0|$. In the sequel we will more in particular assume that $\Lambda = N$ and $\Pi_0 = \{a_i : i \in N\} \cup \{a_\sigma : \sigma \in \Sigma\}$ for L_G.

Falsum ($\bot$), material implication ($\varphi \to \psi$) and disjunction ($\varphi \vee \psi$) are defined as usual. Let further $\langle\pi\rangle\,\varphi$, $\langle i\rangle\,\varphi$, and $\mathtt{while}\ \varphi\ \mathtt{do}\ \pi\ \mathtt{od}$ be short for $\neg[\pi]\neg\varphi$, $\neg\,[i]\,\neg\varphi$, and $(\varphi?;\pi)^*;\neg\varphi?$, respectively.

A frame for a language L is a triple $(S, \{R_\pi\}_{\pi\in\Pi}, \{R_i\}_{i\in\Lambda})$, in which S is a set (of states) and for each $\xi \in \Pi \cup \Lambda$, $R_\xi \subseteq S \times S$ is an *accessibility relation*. A frame is *regular* if the program connectives ";", "$\cup$" and "$*$" have their intuitive interpretations of *sequential composition*, *non-deterministic choice* and *iteration*, respectively. That is, if the following conditions are fulfilled, where $(R_\pi)^*$ denotes the reflexive and transitive closure of R_π:

$$\begin{array}{rcl} R_{\pi_1;\pi_2} & = & R_{\pi_1} \circ R_{\pi_2} \\ R_{\pi_1\cup\pi_2} & = & R_{\pi_1} \cup R\pi_2 \\ R_{\pi^*} & = & (R_\pi)^* \end{array}$$

An *interpretation function* I for a frame $(S, \{R_\pi\}_{\pi\in\Pi}, \{R_i\}_{i\in\Lambda})$ is a function that assigns a subset of S to each propositional variable of the language. A *model* for a language L is a pair consisting of a frame and an interpretation function $I : \Phi_0 \to 2^S$. The semantics of PDL is given by interpreting each formula φ on a model $\mathfrak{M}$ as follows:

$$\begin{array}{lll} \mathfrak{M}, s \models A & \text{iff} & s \in I(A) \\ \mathfrak{M}, s \models \neg\varphi & \text{iff} & \mathfrak{M}, s \not\models \varphi \\ \mathfrak{M}, s \models \varphi \wedge \psi & \text{iff} & \mathfrak{M}, s \models \varphi \text{ and } \mathfrak{M}, s \models \psi \\ \mathfrak{M}, s \models [i]\,\varphi & \text{iff} & \text{for all } s' \in S:\ sR_is' \text{ implies } \mathfrak{M}, s' \models \varphi \\ \mathfrak{M}, s \models [\pi]\,\varphi & \text{iff} & \text{for all } s' \in S:\ sR_\pi s' \text{ implies } \mathfrak{M}, s' \models \varphi \end{array}$$

and setting simultaneously:

$$R_{\varphi?} \quad = \quad \{(s, s') : \mathfrak{M}, s \models \varphi\}.$$

We say that a formula φ is valid in a model $\mathfrak{M}$, $\mathfrak{M} \models \varphi$, if for all states s in the model, $\mathfrak{M}, s \models \varphi$. A formula φ is valid in a frame $\mathfrak{F}$ at a state s, $\mathfrak{F}, s \models \varphi$, if for all interpretation functions I for $\mathfrak{F}$, $(\mathfrak{F}, I)\,, s \models \varphi$. φ is valid in $\mathfrak{F}$ *per se*, $\mathfrak{F} \models \varphi$, if for all interpretation functions I, $(\mathfrak{F}, I) \models \varphi$. Finally, a formula φ of L is valid, $\models \varphi$, if for all frames $\mathfrak{F}$, $\mathfrak{F} \models \varphi$.

4 Games as Frames

Games in extensive form are defined as trees and as such give rise to frames that serve as the semantical entities of our logic. For each extensive game G we define a frame $\mathfrak{F}_G$ for the language L_G as follows:

Definition 5 *Let $G = (V, \prec, Z, M, \iota, u)$ be an extensive game. We define $\mathfrak{F}_G$ as the smallest regular frame $(S, \{R_\pi\}_{\pi\in\Pi}, \{R_i\}_{i\in\Lambda})$ for L_G satisfying the following four requirements:*

$$\begin{array}{lrcl} (1) & S & = & V \\ (2) & sR_is' & \textit{iff} & s \le_i s' \\ (3) & sR_{a_i}s' & \textit{iff} & s \prec s' \textit{ and } \iota(s) = i \\ (4) & sR_{a_\sigma}s' & \textit{iff} & \sigma(s) = s' \end{array}$$

Accordingly, the tree $\left(S, \bigcup_{i\in\Lambda} R_{a_i}\right)$ coincides with $(V, \prec)$. Furthermore, each modality $[i]$ runs over the preference order of player i over the states as induced by the utility function. Hence, $[i]\,\varphi$ intuitively means that φ holds in all states that are at least as preferable to i as the local one. If $sR_{a_i}s'$, this indicates that it is player i's move at s. As a rule, a_i is a non-deterministic program, because a player in a game has several options how to act. The atomic program a_σ is interpreted as the functional relation the strategy profile σ of the game defines on the vertices. Hence, it is inevitably a deterministic program, defined on each of the internal nodes. To illustrate this definition, Figure 2 depicts the frame $\mathfrak{F}_G$ corresponding to the game of example 1 as it is defined for the atomic programs a_{i_1}, a_{i_2} and a_σ for σ such that $\sigma(v_0) = v_2$, $\sigma(v_1) = v_3$, $\sigma(v_2) = z_4$ and $\sigma(v_3) = z_2$.

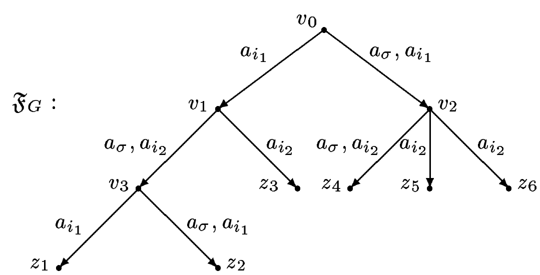

Figure 2. Frame corresponding with Example 1

Since the games that are object of our study have quite a definite tree structure such that each vertex is accessible from the root in a finite number of steps, the class of frames corresponding to a game cannot be fully characterized. The class of frames we are going to consider instead is defined by certain properties that are satisfied by each frame that is induced by a game. These properties are axiomatized by the following schemes:

$$\begin{array}{rc} T_i & [i]\,\varphi \to \varphi \\ 4_i & [i]\,\varphi \to [i]\,[i]\,\varphi \\ D_\sigma & \langle a_\sigma\rangle\,\varphi \to [a_\sigma]\varphi \\ G1_\sigma & \langle a_\sigma\rangle\,\varphi \to \bigvee_{i\in N}\langle a_i\rangle\,\varphi \\ G2_\sigma & \bigvee_{i\in N}\langle a_i\rangle\,\top \to \langle a_\sigma\rangle\,\top \\ G3_i & \langle a_i\rangle\,\top \to \bigwedge_{j\in N\setminus\{i\}}[a_j]\bot \\ G4_{\pi,i} & (\langle\pi\rangle\,\varphi \wedge \langle\pi\rangle\,\psi) \to (\langle\pi\rangle\,(\varphi \wedge \langle i\rangle\,\psi) \vee \langle\pi\rangle\,(\psi \wedge \langle i\rangle\,\varphi)) \end{array}$$

Interpreted on frames corresponding to games these axioms have intuitive readings. Since in each game for each player i the preference relation $\leq_i$ is determined by the utility function u, which range was taken to be the continuum, R_i is a total preorder on the states indeed. Hence, T_i and 4_i, which characterize reflexivity and transitivity of R_i, respectively. The axiom scheme $G4_{\pi,i}$ captures, for each program π, the comparability with respect to R_i of any two states in which π terminates. This fact holds for frames $\mathfrak{F}_G$ because the utility function in each game G is defined for all vertices. $G1$ assures that the strategy profile only prescribes moves the players can perform. Moreover, $G2_\sigma$ makes certain that, whenever a program a_i is enabled so is a_σ, i.e., a player cannot adopt the strategy not to move at the nodes assigned to her. Finally, $G3_i$ guarantees that no two programs a_i and a_j are enabled at one node, i.e., no two players are to play simultaneously. These facts are easily proved and summarized in the following proposition:

Proposition 6 *Let G be a game and $\mathfrak{F} = (S, \{R_\pi\}_{\pi\in\Pi}, \{R_i\}_{i\in\Lambda})$ a frame for L_G (i and j range over Λ, s, s' and s'' over S):*

$$\begin{array}{lll}
\mathfrak{F} \models G1_\sigma & \textit{iff} & R_{a_\sigma} \subseteq \bigcup_{i\in N} R_i \\
\mathfrak{F} \models G2_\sigma & \textit{iff} & \forall s, s', i:\ sR_i s' \textit{ implies } \exists s'' \in S:\ sR_{a_\sigma} s'' \\
\mathfrak{F} \models G3_i & \textit{iff} & \forall s, s', \forall j:\ sR_{a_i} s' \ \&\ sR_{a_j} s' \textit{ imply } i = j \\
\mathfrak{F} \models G4_{\pi,i} & \textit{iff} & \\
\multicolumn{3}{l}{\forall s, s', s'':\ (sR_\pi s' \ \&\ sR_\pi s'') \textit{ imply } (s'R_i s'' \textit{ or } s''R_i s').^4}
\end{array}$$

Let $\mathcal{C}$ be the class of frames F for which $\leq_i$ is reflexive and transitive for all $i \in N$ and which satisfy the conditions on the right-hand side of the equivalences in proposition 6. We have the following fact:

Fact 7 *For all games G:* $\mathfrak{F}_G \in \mathcal{C}$.

5 Characterizing Nash Equilibria

In the previous section we defined for each extensive game G a dynamic language L_G and a frame $\mathfrak{F}_G$ for L_G. Now, some strategy-profiles comprise a Nash-equilibrium in an extensive game, and others do not. If a strategy profile σ is a Nash-equilibrium in an extensive game G, this fact is reflected in certain structural properties of the frame $\mathfrak{F}_G$. What we are after is a formula scheme in L_G for each strategy profile σ, $\vartheta(\sigma)$, such that:

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

First we introduce, for each subset of players $M \subseteq N$ and strategy profile σ, a complex non-deterministic program, $\pi(\sigma, M)$, as an auxiliary notion.

Definition 8 *Consider L_G for some game G. Define for each subset $\{i_0, \ldots, i_k\} \subseteq N$ and for each $\sigma \in \Sigma$ the program $\pi(\sigma, \{i_0, \ldots, i_k\})$ as:*

$$\pi(\sigma, M) \quad := \quad \texttt{while } \langle a_\sigma \rangle \top \texttt{ do } (a_{i_0} \cup \ldots \cup a_{i_k} \cup a_\sigma) \texttt{ od}$$

Usually we will abbreviate $\pi(\sigma, \{a_i\})$ to $\pi(\sigma, a_i)$. Note further that:

$$\pi(\sigma, \varnothing) \quad = \quad \texttt{while } \langle a_\sigma \rangle \top \texttt{ do } a_\sigma \texttt{ od}.$$

The intuition behind this definition becomes clear when we concentrate on frames in the class $\mathcal{C}$. For each $M \subseteq N$, the program $\pi(\sigma, M)$ is non-deterministic if for some $i \in M$, a_i is non-deterministic. $\pi(\sigma, M)$ executes any of the atomic programs a_i with $i \in M$ if it is enabled in a state, the program a_σ otherwise. The program terminates when a_σ is no longer enabled. In any frame satisfying $G2_\sigma$ no $a_i \in \Pi_0$ will be enabled either. In contradistinction, $\pi(\sigma, \varnothing)$ reduces to a deterministic program that repeats a_σ until it terminates.

In any frame $F \in \mathcal{C}$, $R_{a_\sigma} \subseteq R_{\bigcup_{i\in N} a_i}$, and so the program $\pi(\sigma, N)$, when executed in state s, terminates in those states s' that are reachable by a path $s = s_0 R_{a_{i_1}} \ldots R_{a_{i_k}} s_k = s'$ such that for no $j \in N$ and for no state s'', $s' R_{a_j} s''$. The larger the set M, the more non-determinism is brought into the program $\pi(\sigma, M)$, with the deterministic $\pi(\sigma, \varnothing)$ on the one end of the spectrum and $\pi(\sigma, N)$ on the other. In other words, $R_{\pi(\sigma,M)}$ is monotonic in the sense that $M \subseteq M'$ implies $R_{\pi(\sigma,M)} \subseteq R_{\pi(\sigma,M')}$.

For each frame $\mathfrak{F}_G$ and strategy profile σ, the program $\pi(\sigma, N)$ will terminate exactly in the leaf nodes still reachable from the state in which the program is executed. With the program a_σ encoding a strategy profile σ, commencing in s, $\pi(\sigma, \varnothing)$ terminates precisely in that node which σ determines as its unique outcome, i.e., $sR_{\pi(\sigma,\varnothing)}s'$ iff $s' = \hat{\sigma}(s)$. Moreover, as the program a_i is interpreted as the moves available to player i, the possible runs of the program $\pi(\sigma, a_i)$ terminate in exactly the leaf nodes which, by choosing her strategy, i can guarantee the game to end if the other players stick to the strategy profile σ. As such, $\pi(\sigma, i)$ represents the strategy search space of player i the strategies of the other players fixed by σ. In our example, $R_{\pi(\sigma,i_1)}$ is the set $\{(v_0, z_1), (v_0, z_2), (v_0, z_4), (v_1, z_1), (v_1, z_2), (v_2, z_4), (v_3, z_1), (v_3, z_2)\}$, and $\pi(\sigma, i_1)$ is obtained from Figure 2 by removing the leaves z_3, z_5 and z_6, and by substituting a_σ for a_σ, a_{i_2}. Summarizing, the following fact holds.

Fact 9 *Let $G = (V, \prec, Z, N, \iota, u)$ and $\mathfrak{F}_G$ its corresponding frame.*

$$\begin{array}{llll}
(a) & vR_{\pi(\sigma,\varnothing)}z & \textit{iff} & \hat{\sigma}(v) = z \\
(b) & vR_{\pi(\sigma,i)}z & \textit{iff} & \exists \sigma' \in \Sigma\text{: } \sigma \sim_{-i} \sigma' \textit{ and } \hat{\sigma}'(v) = z
\end{array}$$

It is precisely this insight that is exploited in this section to characterize frames for which the strategy program a_σ matches the Nash optimal strategy profile of the corresponding games.

The program $\pi(\sigma, \varnothing)$, as it boils down to an iterated execution of the a_σ program until a final state is reached, combines the strategies of the players as encoded in the strategy profile concerned. Different strategy profiles σ, some of which may be Nash-optimal, will give rise to different $\pi(\sigma, \varnothing)$ programs. The question we address now, is which structural properties a frame $\mathfrak{F}_G$ should comply to, if the program $\pi(\sigma, \varnothing)$ is to mirror a strategy profile that contains a (subgame perfect) best response strategy for a player or one that is a (subgame perfect) Nash-equilibrium. Eventually we will show that these properties can be characterized by a formulae in L_G.

If a player i acts in accordance with her own interest, one would expect her to choose that strategy in her strategy search space which guarantees her the highest payoff. In terms of frames, this can be put as follows. Let a_σ be such that a final state z is reachable by the program $\pi(\sigma, i)$ and another final state z' by $\pi(\sigma, \varnothing)$. Then by monotony $\pi(\sigma, i)$ reaches z' as well. For a_σ to model a strategy profile incorporating an optimal strategy for player i, i should either prefer z to z' or be indifferent between them. Otherwise, i could alter her strategy in such a way that for the resulting strategy profile σ', $\pi(\sigma', \varnothing)$ terminates in z. Formally, this intuitive constraint for a player i to decide at node s and a strategy profile σ can be put as follows:

$$(*) \qquad \text{for all } z, z' \in Z : sR_{\pi(\sigma,\varnothing)}z \text{ and } sR_{\pi(\sigma,i)}z' \text{ imply } z'R_i z.$$

We call $R_{\pi(\sigma,\varnothing)}$ *i-beneficial* in a state v if the property $(*)$ holds for v and i. It turns out that this structural property of the programs

[4] For the proofs of the claims made in this paper the reader be referred to [6], which can also be found at: http://www.cs.uu.nl/people/paulh.

i, a_i and a_σ in a frame $\mathfrak{F}_G$ reflects that σ contains a best-response for i in the game G. By way of illustration, the reader consider once more our example (cf. Figure 2) and the strategy profile σ' such that $\sigma'(v_0) = v_1$, $\sigma'(v_1) = z_3$, $\sigma'(v_3) = z_1$ and $\sigma'(v_2) = z_5$. Then, $R_{\pi(\sigma',\varnothing)}$ is not i_1-beneficial in v_0, for $v_0 R_{\pi(\sigma',\varnothing)} z_3$ and $v_0 R_{\pi(\sigma',i_1)} z_5$ but not $z_5 R_{i_1} z_3$. Neither contains σ' a best-response strategy for i_1 in G. However for the original strategy profile σ, $R_{\pi(\sigma,\varnothing)}$ is i-beneficial in all states for both $i = i_1$ and $i = i_2$ as well as σ is a Nash-equilibrium in G.

In general, the following proposition holds as a result of the facts 7 and 9:

Proposition 10 *Let $G = (V, \prec, Z, N, \iota, u)$ be an extensive game and $\mathfrak{F}_G$ its corresponding frame. Then for $\sigma \in \Sigma$ and $i \in N$ in $\mathfrak{F}_G$:*

- (a) *σ contains a best response for i iff $R_{\pi(\sigma,\varnothing)}$ is i-beneficial in v_0*
- (b) *σ contains a Nash-equilibrium iff for all $i \in N$: $R_{\pi(\sigma,\varnothing)}$ is i-beneficial in v_0*
- (c) *σ contains a s.p. best response for i iff for all $v \in V$: $R_{\pi(\sigma,\varnothing)}$ is i-beneficial in v*
- (b) *σ contains a s.p. Nash-equilibrium iff for all $v \in V$, for all $i \in N$: $R_{\pi(\sigma,\varnothing)}$ is i-beneficial in v*

Now the stage is all set to obtain the central result of this paper. Since for each game G, L_G is, apart from being dynamic, a multi-modal language, we have as a consequence of elementary modal logic the following fact all $\alpha, \beta, \gamma \in \Lambda \cup \Pi$ (cf. [10], pp.64–66):

Fact 11 *For all frames $\mathfrak{F} = (S, \{R_\pi\}_{\pi\in\Pi}, \{R_i\}_{i\in\Lambda})$ for a language L_G. Then for all $\alpha, \beta, \gamma \in \Lambda \cup \Pi$:*

$$\mathfrak{F}, s \models \langle\alpha\rangle [\beta] \varphi \rightarrow [\gamma] \varphi \text{ iff}$$
$$\forall s', s'' \in S: \; sR_\alpha s' \text{ and } sR_\gamma s'' \text{ imply } s' R_\beta s''$$

By choosing $\pi(\sigma, i)$, i and $\pi(\sigma, \varnothing)$ for α, β and γ respectively in the above fact, we procure the following theorem:

Theorem 12 *Let $G = (V, \prec, Z, N, \iota, u)$ be a game and $\mathfrak{F}_G$ be the corresponding frame. Let further $\vartheta(\sigma, i, \varphi)$ abbreviate the formula scheme $\langle\pi(\sigma, i)\rangle [i] \varphi \rightarrow [\pi(\sigma, \varnothing)] \varphi$. For all $i \in N$ and $\sigma \in \Sigma_G$:*

- (a) *σ contains a best response for i iff $\mathfrak{F}_G, v_0 \models \vartheta(\sigma, i, \varphi)$*
- (b) *σ is a Nash-equilibrium iff $\mathfrak{F}_G, v_0 \models \bigwedge_{i\in N} \vartheta(\sigma, i, \varphi)$*
- (c) *σ contains a s.p. best response for i iff $\mathfrak{F}_G \models \vartheta(\sigma, i, \varphi)$*
- (d) *σ is a s.p. Nash-equilibrium iff $\mathfrak{F}_G \models \bigwedge_{i\in N} \vartheta(\sigma, i, \varphi)$*

These results establish that PDL is sufficiently expressible for model checking purposes with a strategic objective. The reservations for us to make in this respect are that the situation concerned can be described as an extensive game of perfect information and that one is interested in the Nash equilibrium solution concept. Alternatively, formula schemes like $\vartheta(\sigma, i, \varphi)$ could be used in a specification of a multi-agent system in which the agents are required to settle on strategies that together constitute a Nash-equilibrium. An interesting question in this respect is whether the program a_σ can be formulated as a complex program that is employed by the players as an algorithm to compute a Nash-optimal choice in each possible circumstance. Still, this issue should be committed to future research.

6 Related and Future Research

This work on incorporating game theory in modal logic, is much inspired by Bonanno's [4]. It differs, however, in three respects. Firstly, Bonanno uses computational tree logic (CTL) rather than PDL. Moreover, his emphasis is on the logical foundations of game-theory rather than the incorporation of game-theoretical notions in logic. Thirdly, his analyses are confined to the notion of backward induction, an algorithm designed to generate subgame perfect Nash-equilibria. Backward induction, however, is only guaranteed to provide a solution in games in which the preferences of the players over the outcomes are *totally* ordered. Independent investigations that are congenial to our approach, are reported by Baltag in [1]. Although his concern is primarily with the epistemic aspects of games, he also proposes a dynamic logic in which Nash-equilibria and related concepts can be characterized. His way of mapping games onto Kripke structures, however, is quite different from ours.

An obvious sequel to the present research would be to incorporate other game-theoretical solution concepts into our dynamic framework. Furthermore, our attention has so far been concentrated on extensive games of perfect information without either repetition or chance moves. In the light of purported applications to the specification of fully-fledged multi-agent systems, this could be considered a major concession. After all, one of the areas where the agent metaphor particularly bears fruit is where the players can only be ascribed partial knowledge of their environment.

These matters merit thorough investigation, as do the intricate epistemic issues of game theory and those related to coalition formation.

ACKNOWLEDGEMENTS

Paul Harrenstein is partly supported by the CABS (Collective Agent Based Systems) project of Delft University of Technology.

REFERENCES

[1] Alexandru Baltag, 'A logic for games', in *Proceedings of the ILLC Workshop on Logic and Games, Held in Amsterdam, November 19-20, 1999*, eds., Marc Pauly and Alexandru Baltag, number PP-1999-25 in ILLC Prepublications Series, pp. 19–20, Amsterdam, (1999). ILLC.

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

[3] Patrick Blackburn, Maarten de Rijke, and Yde Venema, *Modal Logic*, number 53 in Cambridge Tracts in Theoretical Computer Science, Cambridge U.P., Cambridge, 2001.

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

[5] David Harel, Dexter Kozen, and Jerzy Tiuryn, *Dynamic Logic*, MIT Press, Cambridge, Mass. & London, 2001.

[6] Paul Harrenstein, Wiebe van der Hoek, John-Jules Meyer, and Cees Witteveen, 'A modal interpretation of nash-equilibria and some related concepts', in *Proceedings of the Fourth Conference on Logic and the Foundations of Game and Decision Theory (LOFT4), June 30–July 2, 2000*, eds., Bonanno Giacomo, Ernrico Colombatto, and Wiebe van der Hoek, pp. 1–21, Torino, (2000).

[7] John-Jules Meyer, Wiebe van der Hoek, and Bernd van Linder, 'A logical approach to the dynamics of commitments', *Artificial Intelligence*, **113**, 1–40, (1999).

[8] Martin J. Osborne and Ariel Rubinstein, *A Course in Game Theory*, MIT Press, Cambridge, Mass., 1994.

[9] Marc Pauly, *Logic for Social Software*, Ph.D. dissertation, Institute for Logic, Language and Information, Amsterdam, 2001.

[10] Sally Popkorn, *First Steps in Modal Logic*, Cambridge U.P., Cambridge, 1994.

[11] A. S. Rao and M. P. Georgeff, 'Modeling rational agents within a bdi-architecture', in *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning*, eds., J Allen, F Fikes, and E Sandewall, pp. 473–484, San Mateo, CA, (1991). Morgan Kaufmann.