Axioms for Game Logic with Preferences

Sieuwert van Otterloo
Department of Computer Science
University of Liverpool

Wiebe van der Hoek

Michael Wooldridge

Abstract

Game Logic with Preferences (GLP)*, is a logic that makes it possible to reason about how information or assumptions about the preferences of other players can be used by agents in order to realize their own preferences. We extend the work done on this logic by looking at the satisfiability problem for this logic. We introduce an axiom system and show the soundness of this system.*

1. Introduction

We are interested in logics that can be used for the specification, analysis, and verification of mechanisms for social interaction [9]. Examples of such mechanisms include voting and election procedures [4], auction and trading protocols, and solutions for fair division and dispute resolution [2]. These situations can be modeled as extensive games [8, 1], and then one can use game theoretic concepts to characterise properties of the protocol. One would like to model procedures as complete and perfect information games, because this is a simple class of games; but then all agents have complete knowledge of each other's preferences. This is an unrealistic assumption in many situations.

One can also model social procedures as distributed or interpreted systems [5]. These systems are models for the options and knowledge that agents have. Some properties of these systems can be expressed using temporal logics [3]; temporal epistemic logics can be used for knowledge based requirements [5], and coalition (epistemic) logics [10] can be used for expressing abilities of coalitions of agents. Unfortunately, these approaches do not deal with agents' preferences at all. The assumption in these systems is that agents know nothing about other agents' preferences – and of course this is also an unrealistic assumption.

In this paper we use a logic called GLP which has been introduced for reasoning about preferences. This logic is similar to work by Harrenstein et al [6] in that it uses preferences over extensive game forms, but GLP can also look at the consequences of knowing that a coalition of agents act following a preference. It is interpreted over models that do not contain agent's preferences, but assumption on the preferences of agents can be made within formulas. GLP is an extension of propositional modal logic with a set of extra modal operators. These operators, written $[\Gamma : \phi_0]\psi$ express that if it is commonly known that coalition Γ wants to achieve ϕ_0, then ψ holds. In a previous publication [12] we have shown that the model checking problem for this logic has linear computational complexity and presented an implemented model checker[1]. In this paper we extend the work on this logic by looking at the satisfiability problem for this logic.

Section 2 defines and illustrates the logic. In section 3 we present a set of axioms and indicate why these axioms are sound. Section 4 is the conclusion.

2. Definitions

Throughout this paper we follow coalition logic conventions [10] by using Σ to denote a set of agents, Γ for a coalition of agents (so $\Gamma \subseteq \Sigma$), $X, Y \in \Sigma$ as variables over agents, P for a set of propositions and π as an interpretation function for propositions. M is a model, ABC are example agents and ϕ, ψ formulas.

Syntax of GLP We find it convenient to define the syntax of GLP in several stages, beginning with (classical) propositional logic.

Definition 1 *The language PL of propositional logic over a set of propositions P is the smallest set $L \supseteq P$ such that for any $\phi \in L$ and $\psi \in L$ we have that $\phi \vee \phi \in L$ and $\neg\phi \in L$.*

We use propositional logic to express properties of the outcomes or results of games. Propositional logic is *only* used for properties of terminal states, not for intermediate states of a game or protocol. Another important detail is that propositional formulas are not itself GLP formulas.

The logic GLP contains all connectives of propositional logic and two additional operators. One can be used to intro-

1 available at http://www.csc.liv.ac.uk/∼sieuwert/glp

duce formulas of propositional logic in GLP, and the other one is used to express consequences of preferences.

Definition 2 *Let P be a set of propositions, and Σ a group of agents. Let PL be the language of propositional logic over P. The language for* GLP *is the smallest language L such that for any formula $\phi_0 \in PL$ and $\phi, \psi \in L$, it is the case that:*

$$\begin{aligned} \Box\phi_0 &\in L \\ \phi \vee \psi &\in L \\ \neg\phi &\in L \\ [\Gamma : \phi_0]\psi &\in L \end{aligned}$$

The formula $[\Gamma : \phi_0]\psi$ has the intended reading 'In any game in which coalition Γ prefers ϕ_0, ψ will hold'. One can compare this operator with a public announcement operator as in dynamic epistemic logic [11]. The game form represents the uncertainty that there is about the outcome of the game before the game is played. A statement $[\Gamma : \phi_0]\psi$ expresses that ψ holds in the situation after it has become common knowledge that the agents in Γ want ϕ_0 to hold and will act according to this preference. The box operator, $\Box\phi_0$, takes a propositional logic formula ϕ_0, which can be interpreted in a specific state, and converts it into a GLP formula: it means that ϕ_0 holds for every possible outcome. A formula $\Box\phi_0$ thus express that, at the start of the protocol, it is commonly known that the outcome state of the protocol that will be reached satisfies ϕ_0.

A useful shorthand is the sometimes operator $\diamond\phi_0$, which means that ϕ_0 holds for *some* outcome. It is defined as $\diamond\phi_0 = \neg\Box\neg\phi_0$. Where no confusion is possible, we omit brackets and commas in the notation of a set Γ of agents, for example writing $[ABC : \phi_0]\psi$ instead of $[\{A, B, C\} : \phi_0]\psi$. We use ϕ_0, ψ_0 for propositional formulas and ϕ, ψ for GLP formulas to avoid confusion.

INTERPRETATION GLP is interpreted over game form interpretations. These are game trees in which all leaves are annotated with propositions. The assumption is that the game tree is common knowledge between all agents involved, but that the preferences of all agents are initially private. A game form interpretation is a tuple (F, P, π) where F is a game form, P a set of propositions and π an interpretation function that for each leaf s returns the set of propositions that are true in that node. We can interpret propositional formulas over leaves in the following way.

$\pi, s \models p$	iff	$p \in \pi(s)$
$\pi, s \models \phi_0 \vee \psi_0$	iff	$\pi, s \models \phi_0$ or $\pi, s \models \psi_0$
$\pi, s \models \neg\phi_0$	iff	not $\pi, s \models \phi_0$

Formulas of GLP are not interpreted over end states, but over an entire game form. The box operator $\Box\phi_0$ is interpreted in such a way that $M \models \Box\phi_0$ is true if and only if ϕ_0 holds in every leaf of M. Negation, disjunction and conjunction are interpreted in the normal way. The interpretation of $[\Gamma : \phi_0]$ is done using an update function Up. This function takes a coalition Γ, a game form M and a formula ϕ_0, and returns a new game form M'. The new game form M' is a pruned version of M: the agents in Γ have less options to choose from, since they are committed to do actions that guarantee that ϕ_0 holds, if possible. The interpretation of GLP can then be finished with the following definitions.

$M \models \Box\phi_0$	iff	for all $s \in Z(T)$ $\pi, s \models \phi_0$
$M \models \phi \vee \psi$	iff	$M \models \phi$ or $M \models \psi$
$M \models \neg\phi$	iff	not $M \models \phi$
$M \models [\Gamma : \phi_0]\psi$	iff	$Up(M, \Gamma, \phi_0) \models \psi$

We apologize to the reader for not defining formally the update function. Instead we sketch and describe the update function below. A more formal treatment of the interpretation of GLP can be found in [12]. For this publication we have chosen to sketch the intuition behind the update function. First of all the formal treatment is quite long, secondly because this may give more insight. Thirdly many axioms we prove do not heavily depend on the specific update function used. This is for instance the case for distribution, restriction, functionality. The important insight we hope to convey is the general priniciple of using an update function.

To define $Up(M, \Gamma, \phi_0)$ we consider the two player zerosum game in which all the agents in Γ get a payoff of one in outcomes where ϕ_0 holds. The other players in $\Sigma \setminus \Gamma$ get payoff one if ϕ_0 does not hold. These coalitions are thus the two players. Both coalitions want to maximize their payoff.

In two player zerosum games rational behaviour is often related to playing a Nash Equilibrium strategy [8]. Every Nash Equilibrium of a two player zerosum game has the same value, and this value can be computed using backwards induction [8]. The update function is also defined and computed using backwards induction. Let n be a node in the game tree controlled by an agent $X \in \Gamma$. This agent has several options $o_0 \ldots o_N$ to chose from at this node. Some (perhaps none) of these options lead to a subgame in which Γ can get the payoff of one (good moves). In other options, Γ cannot (bad moves). In the updated tree, only the good options are available to X, unless there are only bad options, in which case all bad options are kept. For nodes controlled by agents $Y \notin \Gamma$, all options are kept. In this way the update function $Up(M, \Gamma, \phi_0)$ returns a reduced tree such that the agents of Γ are constrained in their behaviour: they must act toward making ϕ_0 true. A nondeterministic strategy S is more general than S' if they are defined for the same coalition Γ and for all nodes n controlled by Γ it is the case that $S'(n) \subseteq S(n)$. The update function makes no distinction between different good moves or different bad moves.

Therefore the update represents the case were Γ uses the most general 'good' strategy S_Γ.

We distinguish *successful* and *unsuccessful* updates. If $Up(M, \Gamma, \phi_0) \models \Box\phi_0$, then the update is called successful, otherwise unsuccessful. In the original definition of GLP [12], the model $Up(M, \Gamma, \phi_0)$ could have been different from M even in the case of an unsuccessful update. We have found it difficult to specify using axioms the properties of unsuccessful updates, and one line of ongoing research is to investigate whether a slightly different update function might make the problem of finding a sound and complete proof system easier.

Another distinction we can make is that between *abilities* and *side-effects*. A formula of the form $[\Gamma : \phi_0]\psi$ is an ability if $\psi = \Box\phi_0$. Otherwise it is a side-effect. We have found it easier to reason about abilities than about side-effects. For instance the axioms minimax and grand coalition are only about abilities.

Example The example protocol given here is a voting protocol for three agents (A, B, and C) that have to choose between three alternatives (x, y, and z). The requirements are that exactly one alternative is chosen and any group of two agents can force any outcome (majority rules). This is formalized in the next formulas. Let $\Gamma \subseteq \{A, B, C\}$ and u, v variables over the options $O = \{x, y, z\}$.

$\Box(x \vee y \vee z)$	at least one alternative
$\bigwedge_{u,v \in O, u \neq v} \Box\neg(u \wedge v)$	at most one alternative
$\lvert\Gamma\rvert > 1 \Rightarrow \bigwedge_{u \in O}[\Gamma : u]\Box u$	majority decides

In Figure 1, a protocol P_1 is depicted which satisfies these requirements – it is in fact a smallest protocol that satisfies the requirements, as one can verify by testing all smaller trees. It works in two steps. First, A can say whether B or C can make a choice. Then, either B or C can indicate which of the three destinations is chosen. The protocol may seem strange because A cannot directly support a certain outcome. What is the best action for A depends on what B and C will do. In this protocol it is thus important for A to know what the others do, while C and B need no information about the other agents. The next GLP formulas illustrate this peculiarity of protocol P_1.

$$P_1 \models [AB : x]\Box x$$
$$P_1 \models [B : x][A : x]\Box x$$
$$P_1 \not\models [A : x][B : x]\Box x$$

To illustrate the meaning of the update function, the updated model $Up(P_1, \{B\}, x)$ is illustrated in figure 2. One can see that in this model, A can guarantee x by choosing its left option. Therefore, $Up(P_1, \{B\}, x) \models [A : x]\Box x$ and thus $P_1 \models [B : x][A : x]\Box x$.

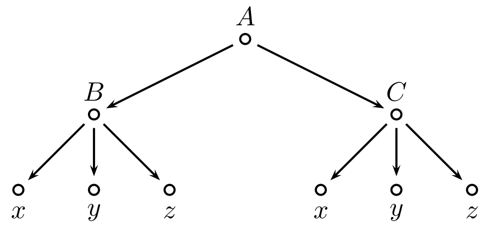

Figure 1. A voting protocol P_1

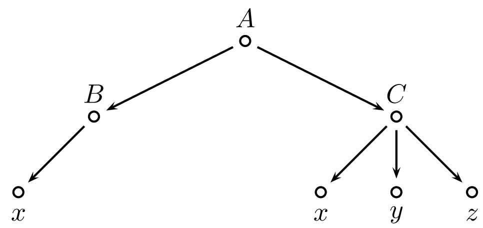

Figure 2. Updated model $Up(P_1, \{B\}, x)$

3. Proof system

In this section we present two reasoning rules and several axioms. The notation used is that $M \models \phi$ means that ϕ holds for model M. The notation $\models \phi$ means that $M \models \phi$ for any model M and in this case we say that ϕ is valid. If one can construct a proof for ϕ using the reasoning rules and axioms below, we say that ϕ is derivable and write $\vdash \phi$. Our major concern in this section is the soundness of the system proposed here. Soundness means that for any formula ϕ we can show that $\vdash \phi$ implies that $\models \phi$.

The system described uses only two reasoning rules: *Modus Ponens* and *Necessitation*. It is not hard to see that these are valid. Since these rules are also present in the system K of modal logic[7], we have omitted these proofs.

$$\frac{\phi \quad \phi \to \psi}{\psi}(MP)$$

$$\frac{\phi}{[\Gamma : \psi_0]\phi}(Nec)$$

The following is a list of axiom schemes. We claim that all these axiom schemes are valid. Most of these axioms are valid. In these schemes, ϕ, χ and ψ stand for any proposition and Γ, Γ_1 for arbitrary coalitions. Σ is the coalition of all agents. τ is an instance of a propositional logic proposition.

1. $\vdash \tau$ tautology
2. $\vdash \Box\tau$ box tautology
3. $\vdash \Box\phi_0 \rightarrow \diamond\phi_0$ box seriality
4. $\vdash \Box(\phi_0 \rightarrow \psi_0) \rightarrow (\Box\phi_0 \rightarrow \Box\psi_0)$ box distribution
5. $\vdash [\Gamma : \chi_0](\phi \rightarrow \psi) \rightarrow ([\Gamma : \chi_0]\phi \rightarrow [\Gamma : \chi_0]\psi)$ distribution
6. $\vdash \Box\phi_0 \rightarrow ([\Gamma : \phi_0]\psi \rightarrow \psi)$ donothing
7. $\vdash [\Gamma : \phi_0]\psi \rightarrow [\Gamma : \phi_0][\Gamma : \phi_0]\psi$ projection
8. $\vdash [\Gamma : \phi_0]\neg\psi \rightarrow \neg[\Gamma : \phi_0]\psi$ functionality
9. $\vdash \Box\phi_0 \rightarrow [\Gamma : \psi_0]\Box\phi_0$ restriction
10. $\vdash \diamond\phi_0 \rightarrow [\Sigma : \Box\phi_0]\Box\phi_0$ grand coalition
11. $\vdash [\Gamma : \phi_0]\Box\phi_0 \leftrightarrow \neg[\Sigma\backslash\Gamma : \neg\phi_0]\Box\neg\phi_0$ minimax
12. $\vdash [\Gamma_1 : \phi_0][\Gamma_2 : \phi_0]\Box\phi_0) \rightarrow [\Gamma_1 \cup \Gamma_2 : \phi_0]\Box\phi_0$ combination
13. $\vdash [\Gamma : \phi_0]\Box\phi_0 \rightarrow [\Gamma_2 \setminus \Gamma : \psi][\Gamma : \phi_0]\Box\phi_0$ preservation
14. $\vdash [\Gamma : \phi_0]\Box\phi_0 \rightarrow [\Gamma : (\phi_0 \vee \psi_0)] \diamond \phi_0$ generality
15. $\vdash [\Gamma : \phi_0]\Box\psi_0 \rightarrow [\Gamma : \psi_0]\Box\psi_0$ effectivity

Proofs of Soundness

$$\models \tau \text{ (tautology)}$$

By definition, a tautology τ evaluates to true under all assignments of truth values to the atoms p, q occurring in the formula. We consider tautologies $\tau(p, q, \ldots)$ where the atoms p, q are replaced by formulas $\Box\phi_0$. The formula with substitutions $\tau(\Box\phi_0, \ldots)$ is thus true regardless of the truth values of $\Box\phi_0, \ldots$. Therefore this axiom is sound.

$$\models \Box\tau \text{ (box tautology)}$$

Because τ is a propositional tautology, it is true in all end nodes of any game form interpretation. Therefore $\Box\tau$ holds in any game form interpretation. This axiom is thus valid.

$$\models \Box\phi_0 \rightarrow \diamond\phi_0 \text{ (seriality)}$$

Let M be any game form interpretation and suppose $M \models \Box\phi_0$. Model M has at least one end state s. Because $M \models \Box\phi_0$ we know that $\pi, s \models \phi_0$. Since there is a state in which ϕ_0 holds, $M \models \diamond\phi_0$.

$$\models \Box(\phi_0 \rightarrow \psi_0) \rightarrow (\Box\phi_0 \rightarrow \Box\psi_0) \text{ (box distribution)}$$

Let M be any game form interpretation and suppose $M \models \Box(\phi_0 \rightarrow \psi_0)$ and $M \models \Box\phi_0$. Let s be any terminal state of M. From the assumptions we know that $\pi, s \models \phi_0$ and $\pi, s \models \phi_0 \rightarrow \psi_0$. This implies that $\pi, s \models \psi_0$. Since we have shown this for an arbitrary end state s, we conclude $M \models \Box\psi_0$.

$$\models [\Gamma : \chi_0](\phi \rightarrow \psi) \rightarrow ([\Gamma : \chi_0]\phi \rightarrow [\Gamma : \chi_0]\psi) \text{ (distribution)}$$

Let M be any game form interpretation and suppose $M \models [\Gamma : \chi_0](\phi \rightarrow \psi)$ and $m \models [\Gamma : \chi_0]\phi$. From the assumptions we know that $Up(M, \Gamma, \chi_0) \models (\phi \rightarrow \psi)$ and $Up(M, \Gamma, \chi_0) \models \phi$. Using modus ponens we derive $Up(M, \Gamma, \chi_0) \models \psi$ and thus we conclude $m \models [\Gamma : \chi_0]\psi$.

$$\models \Box\phi_0 \rightarrow ([\Gamma : \phi_0]\psi \rightarrow \psi) \text{ (donothing)}$$

Let M be any model such that $M \models \Box\phi_0$. The update function removes 'bad moves' of agents $X \in \Gamma$. Since ϕ_0 holds in every possible outcome, there are no 'bad moves' and thus $Up(M, \Gamma, \phi_0) = M$. Therefore $Up(M, \Gamma, \phi_0) \models \psi$ implies that $M \models \psi$.

$$\models [\Gamma : \phi_0]\psi \rightarrow [\Gamma : \phi_0][\Gamma : \phi_0]\psi \text{ (projection)}$$

Let M be any game form interpretation and assume $M \models [\Gamma : \phi_0]\psi$. In that case $Up(M, \Gamma, \phi_0) \models \psi$. In the model $Up(M, \Gamma, \phi_0)$, agents in coalition Γ are doing whatever they can to ensure ϕ_0. Another update with the same formula does not lead to further pruning: $Up(Up(M, \Gamma, \phi_0), \Gamma, \phi_0) = Up(M, \Gamma, \phi_0)$. Therefore $Up(Up(M, \Gamma, \phi_0), \Gamma, \phi_0) \models \psi$ and thus $M \models [\Gamma : \phi_0][\Gamma : \phi_0]\psi$.

$$\models [\Gamma : \phi_0]\neg\psi \leftrightarrow \neg[\Gamma : \phi_0]\psi \text{ (functionality)}$$

Let $M = (\Sigma, N, T, I)$ be any game form interpretation. In the model $Up(M, \Gamma, \phi_0)$, either ψ holds, or $\neg\psi$ holds. In the first case $M \not\models [\Gamma : \phi_0]\neg\psi$ and $M \not\models \neg[\Gamma : \phi_0]\psi$. In the second case $M \models [\Gamma : \phi_0]\neg\psi$ and $M \models \neg[\Gamma : \phi_0]\psi$. The statements on both sides of the double arrow are thus equivalent.

$$\models \Box\phi_0 \rightarrow [\Gamma : \psi_0]\Box\phi_0 \text{ (restriction)}$$

Let M be any game form interpretation and suppose $M \models \Box\phi_0$. This means that in every state s of M we have $\pi, s \models \phi_0$. The set of states S' of the model $Up(M, \Gamma, \psi_0)$ is a subset of the states S of M. Let $s' \in S'$. Since $s' \in S$ we know that $\pi, s \models \phi_0$. Since s' is an arbitrary state of $Up(M, \Gamma, \psi_0)$, we can conclude that $Up(M, \Gamma, \psi_0)\Box\phi_0$.

$$\models \diamond\phi_0 \rightarrow [\Sigma : \Box\phi_0]\Box\phi_0 \text{ (grand coalition)}$$

Let M be any game form interpretation and suppose $M \models \diamond\phi_0$. This means there is a path $n_0 n_1 \ldots n_e$ in the tree associated with the model M from the root to a state n_e where $\pi, n_e \models \phi_0$. Using induction, one can show that at every node n_i the agent that controls node $n_i)$ has an option (namely n_{i+1}) that guarantees ϕ_0. By definition of the update function, this means that in the updated model $Up(M, \Sigma, \phi_0)$, only actions that guarantee ϕ_0 are preserved. Therefore, in every outcome of $Up(M, \Sigma, \phi_0)$, ϕ_0 holds and thus $Up(M, \Sigma, \phi_0) \models \Box\phi_0$.

$$\models [\Gamma : \Box\phi_0]\Box\phi_0 \leftrightarrow \neg[\Sigma\backslash\Gamma : \Box\neg\phi_0]\Box\neg\phi_0 \text{ (minimax)}$$

Let $M = (\Sigma, N, T, I)$ be any game form interpretation. The game in which Γ 'wins' payoff 1 if the outcome satisfies ϕ_0 and $\Sigma\backslash\Gamma$ 'wins' 1 if $\neg\phi_0$ holds in the outcome is finite and thus determined. One of the two coalitions must have a winning strategy. If Γ has a winning strategy, then $[\Gamma : \Box\phi_0]\Box\phi_0$. Otherwise $[\Sigma\backslash\Gamma : \Box\neg\phi_0]\Box\neg\phi_0$.

$$\models [\Gamma_1 : \phi_0][\Gamma_2 : \phi_0]\Box\phi_0 \rightarrow [\Gamma_1 \cup \Gamma_2 : \phi_0]\Box\phi_0 \text{ (combination)}$$

Let M be any model and let $M \models [\Gamma_1 : \phi_0][\Gamma_2 : \phi_0]\Box\phi_0$. This means that $Up(Up(M, \Gamma_1, \phi_0), \Gamma_1, \phi_0) \models \Box\phi_0$. So there is a nondeterministic strategy for Γ_1 and one for Γ_2 such that when both are adhered to, an outcome that makes ϕ true is guaranteed. These two strategies can be combined into one strategy for $\Gamma_1 \cup \Gamma_2$, and therefore this larger coalition has a strategy for ensuring ϕ. Thus $[\Gamma_1 \cup \Gamma_2 : \phi_0]\Box\phi_0$.

$$\models [\Gamma : \phi_0]\Box\phi_0 \rightarrow [\Gamma_2 \setminus \Gamma : \psi_0][\Gamma : \phi_0]\Box\phi_0 \text{ (preservation)}$$

Let M be any game form interpretation and assume $M \models [\Gamma : \phi_0]\Box\phi_0$. This means that Γ has a strategy S such that using strategy S ensures an outcome satisfying $\Box\phi_0$. In the model after the update $Up(M, \Gamma_2 \setminus \Gamma, \psi_0)$ this strategy can still be applied, and will still guarantee ϕ_0. In this model, Γ can thus achieve ϕ_0 and therefore $M \models [\Gamma_2 \setminus \Gamma : \psi_0][\Gamma : \phi_0]\Box\phi_0$.

$$\models [\Gamma : \phi_0]\Box\phi_0 \rightarrow [\Gamma : (\phi_0 \vee \psi_0)] \diamond \phi_0 \text{ (generality)}$$

Let M be any game form interpretation and assume $M \models [\Gamma : \phi_0]\Box\phi_0$. This means that Γ has a strategy S that ensures ϕ_0. This strategy makes $(\phi_0 \vee \psi_0)$ hold. In the updated model $M' = Up(M, \Gamma, \phi_0 \vee \psi_0)$ the coalition Γ plays a most general strategy and thus a strategy that is as least as general than S. Therefore the outcomes that were part of $Up(M, \Gamma_1, \phi_0)$ are still present in M' and thus $M' \models \diamond\psi_0$.

$$\vdash [\Gamma : \phi_0]\Box\psi_0 \rightarrow [\Gamma : \psi_0]\Box\psi_0 \text{ (effectivity)}$$

Let M be any game form interpretation and assume $M \models [\Gamma : \phi_0]\Box\psi_0$. This means that Γ has a strategy S that ensures ψ_0. Which means $M \models [\Gamma : \psi_0]\Box\psi_0$.

4. Conclusion

This paper looks into the problem of finding an axiomatization for the logic GLP. This is still work in progress. We have presented a system of axioms that is sound. We have not proven completeness of this system but hope to do that in the future.

References

[1] K. Binmore. *Fun and Games: A Text on Game Theory*. D. C. Heath and Company: Lexington, MA, 1992.

[2] S Brahms and D Taylor. *Fair division: from cake cutting to dispute resolution*. Cambridge University Press, 1996.

[3] E. M. Clarke, O. Grumberg, and D. A. Peled. *Model Checking*. The MIT Press: Cambridge, MA, 2000.

[4] V. Conitzer and T. Sandholm. Complexity of mechanism design. In *Proceedings of the Uncertainty in Artificial Intelligence Conference (UAI), Edmonton, Canada.*, 2002.

[5] R. Fagin, J. Halpern, Y. Moses, and M. Vardi. *Reasoning about knowledge*. The MIT Press: Cambridge, MA, 1995.

[6] B. Harrenstein, W. van der Hoek, J.-J. Ch. Meyer, and C. Witteveen. On modal logic interpretations of games. In *Procs ECAI 2002*, volume 77, pages 28–32, Amsterdam, July 2002.

[7] J.-J. Ch. Meyer and W. van der Hoek. *Epistemic Logic for AI and Computer Science*. Cambridge University Press: Cambridge, England, 1995.

[8] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. The MIT Press: Cambridge, MA, 1994.

[9] M. Pauly. *Logic for Social Software*. PhD thesis, University of Amsterdam, 2001. ILLC Dissertation Series 2001-10.

[10] W. van der Hoek and M. Wooldridge. Cooperation, knowledge, and time: Alternating-time temporal epistemic logic and its applications. *Studia Logica*, 75(4):125–157, 2003.

[11] H. P. van Ditmarsch. *Knowledge Games*. PhD thesis, University of Groningen, Groningen, 2000.

[12] S. van Otterloo, W. van der Hoek, and M. Wooldridge. Preferences in game logics. In *AAMAS 2004*, New York, July 2004.