

# Possibilistic Boolean Games: Strategic Reasoning under Incomplete Information

Sofie De Clercq<sup>1</sup>, Steven Schockaert<sup>2</sup>, Martine De Cock<sup>1,3</sup>, and Ann Nowé<sup>4</sup>

<sup>1</sup> Dept. of Applied Math., CS & Stats, Ghent University, Ghent, Belgium  
{SofieR.DeClercq, Martine.DeCock}@ugent.be

<sup>2</sup> School of Computer Science & Informatics, Cardiff University, Cardiff, UK  
S.Schockaert@cs.cardiff.ac.uk

<sup>3</sup> Center for Data Science, University of Washington Tacoma, USA  
MDeCock@uw.edu

<sup>4</sup> Computational Modeling Lab, Vrije Universiteit Brussel, Brussels, Belgium  
ANowe@vub.ac.be

**The final publication is available at Springer via  
[http://dx.doi.org/10.1007/978-3-319-11558-0\\_14](http://dx.doi.org/10.1007/978-3-319-11558-0_14)**

# Possibilistic Boolean Games: Strategic Reasoning under Incomplete Information

Sofie De Clercq<sup>1</sup>, Steven Schockaert<sup>2</sup>, Martine De Cock<sup>1,3</sup>, and Ann Nowé<sup>4</sup>

<sup>1</sup> Dept. of Applied Math., CS & Stats, Ghent University, Ghent, Belgium  
{SofieR.DeClercq, Martine.DeCock}@ugent.be

<sup>2</sup> School of Computer Science & Informatics, Cardiff University, Cardiff, UK  
S.Schockaert@cs.cardiff.ac.uk

<sup>3</sup> Center for Data Science, University of Washington Tacoma, USA  
MDeCock@uw.edu

<sup>4</sup> Computational Modeling Lab, Vrije Universiteit Brussel, Brussels, Belgium  
ANowe@vub.ac.be

**Abstract.** Boolean games offer a compact alternative to normal-form games, by encoding the goal of each agent as a propositional formula. In this paper, we show how this framework can be naturally extended to model situations in which agents are uncertain about other agents' goals. We first use uncertainty measures from possibility theory to semantically define (solution concepts to) Boolean games with incomplete information. Then we present a syntactic characterization of these semantics, which can readily be implemented, and we characterize the computational complexity.

## 1 Introduction

Boolean games (BGs) are games in which the agents' goals are formalized using propositional formulas [12]. The atomic propositional variables occurring in these goals are called the action variables, since each of them is controlled by one agent. Originally, BGs were introduced with binary preferences, i.e. the goal of an agent is a single propositional formula and the utility of an agent is 1 if its goal is satisfied and 0 otherwise [12]. Various suggestions have been made in the literature to overcome this limitation of expressiveness. One approach is the introduction of costs on the action variables [10]. Another suggestion is a generalization of the BG framework towards compact preference relations on the set of outcomes, e.g. by using a prioritized goal base per agent [4]. Recently, the limitation has also been overcome by replacing the classical two-valued logic for representing the goals by many-valued Łukasiewicz logic [13]. This extension allows many degrees to which a goal can be satisfied, as opposed to the sole distinction between satisfaction or non-satisfaction. In this paper, we consider a variant of BGs with prioritized goal bases. An agent is most eager to achieve the goal with the highest priority. If this goal cannot be achieved, the agent will settle with the satisfaction of the goal with the second-highest priority, etc.

*Example 1.* Bob and Alice are going out. Alice – agent 1 – controls action variable  $a$ , and Bob – agent 2 – controls  $b$ . Setting their action variable to true corresponds to attending a sports game; setting it to false corresponds to going to the theatre. Bob and

Alice’s first priority is to go out together. If they do not go out together, Bob prefers a sports game, whereas Alice prefers the theatre. This can be represented with a preference ordering over the outcomes per agent or with a pay-off matrix:

$$\begin{array}{l}
 (a, b) =_1 (\neg a, \neg b) >_1 (\neg a, b) >_1 (a, \neg b) \\
 (a, b) =_2 (\neg a, \neg b) >_2 (\neg a, b) >_2 (a, \neg b)
 \end{array}
 \quad
 \begin{array}{c|c|c}
 \text{Bob} \setminus \text{Alice} & a & \neg a \\
 \hline
 b & (2, 2) & (1, 1) \\
 \hline
 \neg b & (0, 0) & (2, 2)
 \end{array}$$

There are 2 pure Nash equilibria – outcomes such that no one has an incentive to deviate: attending a sports game together and going to the theatre together.

Our aim in this paper is to propose an extension to the BG framework in which agents can be uncertain about other agents’ goals. An important concern is that the resulting framework should still enable a compact and intuitive representation of games, as these are the main strengths of BGs. We therefore introduce a compact syntactic framework, which we prove to correspond to an intuitive semantic framework. Using our extended BG framework, we aim to determine rational behaviour for agents which are uncertain about the other agents’ goals.

Although uncertainty in game theory has been studied extensively (see e.g. [14]), the literature on BGs with incomplete information is currently limited. Uncertainty can be either epistemic or stochastic of nature. The former is caused by incomplete knowledge about the game, whereas the latter is e.g. caused by actions which do not always have the same effect on the outcome. This paper concerns epistemic uncertainty. To the best of our knowledge, the existing work on BGs with uncertainty also falls in the category of epistemic uncertainty. However, in contrast to our work, the uncertainty is not related to the goals. Grant et al. [11] incorporate uncertainty in the BG framework by introducing a set of environment variables outside the control of any agent. Each agent has some (possibly incorrect) belief about the value of the environment variables. The focus of [11] is to manipulate the BGs by making announcements about the true value of some environment variables, in order to create a stable solution if there were none without the announcements. Ågotnes et al. [2] address uncertainty in BGs by extending the framework of BGs with a set of observable action variables for every agent, i.e. every agent can only observe the values assigned to a particular subset of action variables. As a result, agents are not able to distinguish between some strategy profiles, if these profiles only differ in action variables that are not observable to that agent. Three notions of verifiable equilibria are investigated, capturing respectively strategy profiles for which all agents know that they *might be* pure Nash equilibria (PNEs), strategy profiles for which all agents know that they *are* PNEs and strategy profiles for which it is common knowledge that they are PNEs, i.e. all agents know that they are PNEs and all agents know that all agents know that they are PNEs etc. The same authors have extended this framework to epistemic BGs [1], in which the logical language for describing goals is broadened to a multi-agent epistemic modal logic. Note, however, that agents are still completely aware of each others’ goals in this framework.

In this paper, we study BGs with incomplete information, considering agents which have their own beliefs about the goals of other agents. Although probability theory is often used to model uncertainty in game theory [14], a possibilistic logic approach provides a simple and elegant mechanism for modeling partial ignorance, which is closely related to the notion of epistemic entrenchment [8]. Being based on ranking formulas

(at the syntactic level) or possible worlds (at the semantic level), possibilistic logic has the advantage of staying close to classical logic. As a result, we will be able to introduce methods for solving possibilistic BGs that are entirely similar to methods for solving standard BGs.

*Example 2.* Consider again the scenario of Example 1, but assume that Bob and Alice are unaware of each other’s goals. If Bob’s knowledge of Alice’s goal is correct, but Alice thinks that Bob does not want to join her to the theatre, then, based on their beliefs, attending a sports game together is a ‘better’ solution than going to the theatre together. Indeed, Alice believes that Bob will not agree to go to the theatre together, but they both believe that the other will agree to attend a sports game together.

The paper is structured as follows. First, we briefly recall possibility theory and BGs. In Section 3, we construct the framework of BGs with uncertainty, both from an intuitive semantic and a compact syntactic point of view. Moreover, we show that the proposed semantic and syntactic definitions are equivalent, and we characterise the complexity of the associated decision problems.

## 2 Preliminaries

In this section, we recall possibilistic logic and Boolean games. As usual, the logical language  $L_\Phi$  associated with a finite set of atomic propositional variables (atoms)  $\Phi$  contains the following formulas:

- every propositional variable of  $\Phi$ ,
- the logical constants  $\perp$  and  $\top$ , and
- the formulas  $\neg\varphi$ ,  $\varphi \rightarrow \psi$ ,  $\varphi \leftrightarrow \psi$ ,  $\varphi \wedge \psi$  and  $\varphi \vee \psi$  for every  $\varphi, \psi \in L_\Phi$ .

An interpretation of  $\Phi$  is defined as a subset  $\omega$  of  $\Phi$ , with the convention that all atoms in  $\omega$  are interpreted as true ( $\top$ ) and all atoms in  $\Phi \setminus \omega$  are interpreted as false ( $\perp$ ). An interpretation can be extended to  $L_\Phi$  in the usual way. If a formula  $\varphi \in L_\Phi$  is true in an interpretation  $\omega$ , we denote this as  $\omega \models \varphi$ .

### 2.1 Possibilistic Logic

Possibilistic logic (see e.g. [9] for a more comprehensive overview) is a popular tool to encode and reason about uncertain information in an intuitive and compact way.

**Definition 1 (Possibility Distribution).** *Let  $\Omega$  be a finite universe. A possibility distribution on  $\Omega$  is a mapping  $\pi : \Omega \rightarrow [0, 1]$ .*

In possibilistic logic, given a logical language  $L_\Phi$ , the set of interpretations of  $\Phi$  is used as the universe of a possibility distribution. If  $\pi(\omega) = 1$ ,  $\omega$  is considered to be completely possible, whereas  $\pi(\omega) = 0$  corresponds to  $\omega$  being completely impossible. Available information encodes which worlds cannot be excluded based on available knowledge. Therefore, smaller possibility degrees are more specific, as they rule out more possible worlds. A possibility distribution such that  $\pi(\omega) = 1$  for every  $\omega \in \Omega$

thus corresponds to a state of complete ignorance. Note that a possibility distribution is not the same as a probability distribution, since we do not require that  $\sum_{\omega \in \Omega} \pi(\omega) = 1$ . An ordering  $\leq$  on all possibility distributions on  $\Omega$  can be defined as  $\pi_1 \leq \pi_2$  iff it holds that  $\pi_1(\omega) \leq \pi_2(\omega)$ ,  $\forall \omega \in \Omega$ , assuming the natural ordering on  $[0, 1]$ . We say that  $\pi_1$  is at least as specific as  $\pi_2$  when  $\pi_1 \leq \pi_2$ . The maximal elements w.r.t.  $\leq$  are called the least specific possibility distributions. A possibility and necessity measure are induced by a possibility distribution in the following way.

**Definition 2 (Possibility and Necessity Measure).** *Given a possibility distribution  $\pi$  in a universe  $\Omega$ , the possibility  $\Pi(A)$  and necessity  $N(A)$  that an event  $A \subseteq \Omega$  occurs is defined as:*

$$\Pi(A) = \sup_{\omega \in A} \pi(\omega); \quad N(A) = \inf_{\omega \notin A} (1 - \pi(\omega))$$

In possibilistic logic, we abbreviate  $N(\{\omega \in \Omega \mid \omega \models \varphi\})$  as  $N(\varphi)$  for a formula  $\varphi$ .

**Definition 3 (Possibilistic Knowledge Base).** *Let  $\Phi$  be a set of atoms. A finite set  $\{(\varphi_1, \alpha_1), \dots, (\varphi_m, \alpha_m)\}$  of pairs of the form  $(\varphi_i, \alpha_i)$ , with  $\varphi_i \in L_\Phi$  and  $\alpha_i \in ]0, 1]$ , is a possibilistic knowledge base (KB). It encodes a possibility distribution, namely the least specific possibility distribution satisfying the constraints  $N(\varphi_i) \geq \alpha_i$ .*

The possibility distribution  $\pi_{\mathcal{K}}$  encoded by a KB  $\mathcal{K}$  is well-defined because there is a unique least specific possibility distribution which satisfies the constraints of  $\mathcal{K}$  [7].

The necessity measure  $N$  satisfies the property  $N(p \wedge q) = \min(N(p), N(q))$ . The following inference rules are associated with possibilistic logic:

- $(\neg p \vee q, \alpha); (p \vee r, \beta) \vdash (q \vee r, \min(\alpha, \beta))$  (resolution rule),
- if  $p$  entails  $q$  classically, then  $(p, \alpha) \vdash (q, \alpha)$  (formula weakening),
- for  $\beta \leq \alpha$ ,  $(p, \alpha) \vdash (p, \beta)$  (weight weakening),
- $(p, \alpha); (p, \beta) \vdash (p, \max(\alpha, \beta))$  (weight fusion).

The axioms consist of all propositional axioms with weight 1. These inference rules and axioms are sound and complete in the following sense [7]: it holds that  $\mathcal{K} \vdash (\varphi, \alpha)$  iff  $N(\varphi) \geq \alpha$  for the necessity measure  $N$  induced by  $\pi_{\mathcal{K}}$ . Another useful property is  $\mathcal{K} \vdash (\varphi, \alpha)$  iff  $\mathcal{K}_\alpha \vdash \varphi$  (in the classical sense) [9], with  $\mathcal{K}_\alpha = \{\varphi \mid (\varphi, \beta) \in \mathcal{K}, \beta \geq \alpha\}$  the  $\alpha$ -cut of  $\mathcal{K}$ .

## 2.2 Boolean Games

We use a generalization of the notion of Boolean games [5] by allowing agents to have non-dichotomous utilities. This approach is a variant of the BGs with prioritized goal bases considered in [4]. Our notation is based on [2].

**Definition 4 (Boolean Game [4]).** *A Boolean game (BG) is a tuple  $G = (\Phi_1, \dots, \Phi_n, \Gamma_1, \dots, \Gamma_n)$ . The set of agents  $\{1, \dots, n\}$  is denoted as  $N$ . For every  $i \in N$ ,  $\Phi_i$  is a finite set of propositional variables, disjoint with  $\Phi_j, \forall j \neq i$ . We denote  $\Phi = \bigcup_{i \in N} \Phi_i$ . For every  $i \in N$ ,  $\Gamma_i = \{\gamma_i^1; \dots; \gamma_i^p\}$  is  $i$ 's prioritized goal base. The formula  $\gamma_i^j \in L_\Phi$  is agent  $i$ 's goal of priority  $j$ . We assume that the number of priority levels  $p$  is fixed for all agents.*

The set  $\Phi$  contains all action variables. Agent  $i$  can set the variables under its control, i.e. those in  $\Phi_i$ , to true or false. Note that every variable is controlled by exactly one agent. By convention, priority numbers are ordered from high priority (level 1) to low priority (level  $p$ ). Definition 4 corresponds to a particular case of generalized BGs [4], in which the preference relation is total for every agent. The results presented in this paper can easily be generalized to accommodate for partially ordered preference relations. However, as modeling preferences is not the focus of this paper, we prefer the simpler setting of Definition 4, for clarity.

**Definition 5 (Strategy Profile [2]).** For each agent  $i \in N$ , an interpretation of  $\Phi_i$  is called a strategy of  $i$ . An  $n$ -tuple  $\nu = (\nu_1, \dots, \nu_n)$ , with  $\nu_i$  a strategy of agent  $i$ , is called a strategy profile or outcome of  $G$ .

Because  $\{\Phi_1, \dots, \Phi_n\}$  is a partition of  $\Phi$  and  $\nu_i \subseteq \Phi_i, \forall i \in N$ , we also (ab)use the set notation  $\bigcup_{i \in N} \nu_i \subseteq \Phi$  for a strategy profile  $\nu = (\nu_1, \dots, \nu_n)$ . We denote the set of all strategy profiles as  $\mathcal{V}$ . With  $\nu_{-i}$  we denote the projection of the strategy profile  $\nu = (\nu_1, \dots, \nu_n)$  on  $\Phi_{-i} = \Phi \setminus \Phi_i$ , i.e.  $\nu_{-i} = (\nu_1, \dots, \nu_{i-1}, \nu_{i+1}, \dots, \nu_n)$ . If  $\nu'_i$  is a strategy of agent  $i$ , then  $(\nu_{-i}, \nu'_i)$  is a shorthand for  $(\nu_1, \dots, \nu_{i-1}, \nu'_i, \nu_{i+1}, \dots, \nu_n)$ .

The utility for every agent  $i$  follows naturally from the satisfaction of its goals.

**Definition 6 (Utility Function).** For each  $i \in N$  and  $\nu \in \mathcal{V}$ , the utility for  $\nu$  is defined as  $u_i(\nu) = p + 1 - \min\{k \mid 1 \leq k \leq p, \nu \models \gamma_i^k\}$ , with  $\min \emptyset = p + 1$  by convention.

Note that the specific utility values do not matter since the solution concepts that we will discuss in this paper are qualitative; only the preference ordering  $\geq_i$  on  $\mathcal{V}$  induced by the utility function  $u_i$  is relevant:  $\nu \geq_i \nu'$  iff  $u_i(\nu) \geq u_i(\nu'), \forall \nu, \nu' \in \mathcal{V}$ . A common qualitative solution concept in game theory is the notion of pure Nash equilibrium.

**Definition 7 (Pure Nash Equilibrium).** A strategy profile  $\nu = (\nu_1, \dots, \nu_n)$  for a BG  $G$  is a pure Nash equilibrium (PNE) iff for every agent  $i \in N$ ,  $\nu_i$  is a best response (BR) to  $\nu_{-i}$ , i.e.  $u_i(\nu) \geq u_i(\nu_{-i}, \nu'_i), \forall \nu'_i \subseteq \Phi_i$ .

*Example 1 (continued).* Recall the scenario of Example 1. Alice and Bob's goal bases can be written as  $\Gamma_1 = \{a \leftrightarrow b; \neg a\}$  and  $\Gamma_2 = \{a \leftrightarrow b; b\}$ . This encoding naturally captures the fact that e.g. Bob's first priority is to go out with Alice and his second priority is to attend a sports game. Both agents have utility 2 in the PNEs  $\{a, b\}$  and  $\emptyset$ .

## 3 Boolean Games with Incomplete Information

### 3.1 Semantic Approach

Consider a set of agents  $N$ , controlling the action variables in  $\Phi_1, \dots, \Phi_n$ , who are uncertain about each other's goals. Let us denote the set of possible goal bases with  $p$  levels as  $\mathcal{G} = \{\{\gamma^1; \dots; \gamma^p\} \mid \forall k \in \{1, \dots, p\} : \gamma^k \in L_\Phi \text{ in conjunctive normal form and } (k \neq p \Rightarrow \gamma^k \models \gamma^{k+1})\}$ . Note that any formula can be transformed into an equivalent formula in conjunctive normal form (CNF) and that any goal base  $\{\gamma^1; \dots; \gamma^p\}$  violating the condition  $\gamma^k \models \gamma^{k+1}, \forall k \neq p$  can be transformed into a semantically equivalent goal base which does satisfy the property, namely  $\{\gamma^1; \gamma^1 \vee$

$\gamma^2; \dots; \bigvee_{m=1}^p \gamma^m$ . Moreover, all agents have the same set of possible goal bases. Let us define  $\mathcal{BG}(\Phi_1, \dots, \Phi_n) = \{(\Phi_1, \dots, \Phi_n, \Gamma_1, \dots, \Gamma_n) \mid \Gamma_1, \dots, \Gamma_n \in \mathcal{G}\}$  as the set of all possible BGs, given the considered partition of action variables. When the partition  $\Phi_1, \dots, \Phi_n$  is clear from the context, we abbreviate  $\mathcal{BG}(\Phi_1, \dots, \Phi_n)$  as  $\mathcal{BG}$ . The knowledge of an agent  $i$  about the goals of the other agents can be captured by a possibility distribution  $\pi_i$  over  $\mathcal{BG}$ , encoding  $i$ 's beliefs about what is the actual game being played.

*Example 2.* Recall the scenario of Example 1. Suppose Bob has perfect knowledge of Alice's preferences, then  $\pi_2 : \mathcal{BG} \rightarrow \{0, 1\}$  maps every BG to 0, except the BGs with the preference orderings of Example 1, i.e. the actual game being played is the only one considered possible by Bob. Suppose Alice is certain that Bob wants to attend a sports game together, or attend the game on his own if attending it together is not possible. Then  $\pi_1 : \mathcal{BG} \rightarrow \{0, 1\}$  maps all BGs to 0, except those with the preference orderings

$$\begin{array}{l} \{a, b\} =_1 \emptyset >_1 \{b\} >_1 \{a\} \\ \{a, b\} >_2 \{b\} >_2 \emptyset =_2 \{a\} \end{array} \quad \begin{array}{c|c|c} \text{Bob} \setminus \text{Alice} & a & \neg a \\ \hline b & (2, 2) & (1, 1) \\ \hline \neg b & (0, 0) & (0, 2) \end{array}$$

Our first aim is to determine to which degree a specific strategy profile  $\nu$  is necessarily/possibly a PNE according to agent  $i$ . Intuitively, it is possible to degree  $\lambda$  that a strategy profile  $\nu$  is a PNE according to  $i$  iff there exists a BG  $G \in \mathcal{BG}$  such that  $\nu$  is a PNE in  $G$  and such that  $i$  considers it possible to degree  $\lambda$  that  $G$  is the real game being played, i.e.

$$\Pi_i(\{G \in \mathcal{BG} \mid \nu \text{ is a PNE in } G\}) = \lambda$$

Similarly, it is certain to degree  $\lambda$  that a strategy profile  $\nu$  is a PNE according to  $i$  iff for every  $G \in \mathcal{BG}$  such that  $\nu$  is no PNE, it holds that  $i$  considers it possible to degree at most  $1 - \lambda$  that  $G$  is the real game being played, i.e.

$$N_i(\{G \in \mathcal{BG} \mid \nu \text{ is a PNE in } G\}) = \lambda$$

Using the previously introduced degrees, we can define measures which offer a way to distinguish between multiple equilibria, motivated by Schellings' notion of focal points [15]. An equilibrium is a focal point if, for some reason other than its utility, it stands out from the other equilibria. In our case, the reason can be that agents have a higher certainty that the outcome is actually a PNE, using the degrees to which a strategy profile is necessarily a PNE. Note that there might not exist an outcome which every agent believes is necessarily a PNE, even when the (unknown) game being played has one or more PNEs. In such cases, the degree to which various strategy profiles are possibly a PNE could be used to guide decisions.

**Definition 8.** Given the possibility measures  $\Pi_i$  for every  $i$ , the degree to which all agents find it possible that the strategy profile  $\nu$  is a PNE is

$$\text{poss}(\nu) = \min_{i \in N} \Pi_i(\{G \in \mathcal{BG} \mid \nu \text{ is a PNE in } G\})$$

Similarly, given the necessity measures  $N_i$  for every  $i$ , the degree to which all agents find it necessary that  $\nu$  is a PNE is defined as

$$\text{nec}(\nu) = \min_{i \in N} N_i(\{G \in \mathcal{BG} \mid \nu \text{ is a PNE in } G\})$$

### 3.2 Syntactic Approach

While the concepts from Section 3.1 define useful notions w.r.t. the possibility or necessity that agents play best responses or that strategy profiles are PNEs, they cannot be applied in practice, since the number of BGs in  $\mathcal{BG}$  is exponential. In this section, we present a syntactic counterpart which will allow for a more compact representation of the agents knowledge about the game being played.

**Definition 9 (Goal-Knowledge Base).** Agent  $i$ 's knowledge about the goals of agent  $j$  is encoded in a goal knowledge base  $\mathcal{K}_i^j$  of  $i$  w.r.t.  $j$  containing formulas of the form  $(\varphi \rightarrow g_j^k, \lambda)$ ,  $(\varphi \leftarrow g_j^k, \lambda)$  or  $(\varphi \leftrightarrow g_j^k, \lambda)$ , where  $1 \leq k \leq p$ ,  $\varphi \in L_\Phi$ ,  $\lambda \in ]0, 1]$  and  $g_j^k$  a new atom, encoding  $j$ 's goal of priority  $k$ . A goal-KB  $\mathcal{K}_i^j$  is goal-consistent, i.e. for every  $\varphi, \psi \in L_\Phi$  such that  $(\varphi \rightarrow g_j^k, \lambda) \in \mathcal{K}_i^j$  and  $(\psi \leftarrow g_j^k, \lambda) \in \mathcal{K}_i^j$ , it holds that  $\varphi \models \psi$  classically. Moreover,  $\mathcal{K}_i^j$  contains  $\{(g_j^k \rightarrow g_j^{k+1}, 1) \mid 1 \leq k \leq p-1\}$ .

A goal-KB  $\mathcal{K}_i^j$  captures the knowledge of agent  $i$  about the goal base of agent  $j$ . In our examples, the formulas  $\{(g_j^k \rightarrow g_j^{k+1}, 1) \mid 1 \leq k \leq p-1\}$ , which belong to  $\mathcal{K}_i^j$  by definition, are not explicitly mentioned. These formulas express that, if agent  $j$ 's utility is at least  $p+1-k$ , it is at least  $p-k$ . Furthermore, the information that we like to express in  $\mathcal{K}_i^j$  exists of necessary and/or sufficient conditions for the utility of agent  $j$ . For instance, agent  $i$  might believe that with certainty  $\lambda$ ,  $\varphi$  is a sufficient condition for satisfying the goal with priority  $k$ , i.e. achieving a utility of at least  $p+1-k$ . This is encoded as  $(\varphi \rightarrow g_j^k, \lambda) \in \mathcal{K}_i^j$ . Similarly, agent  $i$  might believe with certainty  $\lambda$  that  $\varphi$  is a necessary condition for achieving the goal with priority  $k$ , i.e.  $(\varphi \leftarrow g_j^k, \lambda) \in \mathcal{K}_i^j$ . These types can be combined as  $(\varphi \leftrightarrow g_j^k, \lambda) \in \mathcal{K}_i^j$ . Note how adding the atoms  $g_j^k$  to the language allows us to explicitly encode what an agent knows about the goal of another agent. This is inspired by the approach from [16] for merging conflicting sources, where similarly additional atoms are introduced to encode knowledge about the unknown meaning of vague properties, in the form of necessary and sufficient conditions.

*Example 4.* Recall the scenario of Example 1. Suppose Bob has a good idea of what Alice's goal base looks like:  $\mathcal{K}_2^1 = \{((a \leftrightarrow b) \leftrightarrow g_1^1, 0.9), (((a \leftrightarrow b) \vee \neg a) \leftrightarrow g_1^2, 0.6)\}$ . He is very certain that Alice's first priority is to go out together and rather certain that she prefers the theatre in case they do not go out together. Although Alice is very certain that Bob will be pleased if they attend a sports game together, she is only a little certain whether Bob would be just as pleased if they attend the cultural event together. She knows Bob prefers to go a sports game as a second priority. Her knowledge of Bob's goal base can be captured by  $\mathcal{K}_1^2 = \{((a \wedge b) \rightarrow g_2^1, 0.8), ((\neg a \wedge \neg b) \rightarrow g_2^1, 0.3), (b \rightarrow g_2^2, 1)\}$ .

It is natural to assume that  $\mathcal{K}_i^i = \{(g_i^k \leftrightarrow \bigvee_{m=1}^k \gamma_i^m, 1) \mid k \in \{1, \dots, p\}\}$ , i.e. every agent knows its own goal base and the corresponding utility. However, this assumption is not necessary for the results in this paper. By requiring goal-consistency, we ensure that the knowledge base  $\mathcal{K}_i^j$  only encodes beliefs about the goal of agent  $j$ . Without this assumption, it could be possible to derive from  $\mathcal{K}_i^j$  formulas of the form  $\varphi \rightarrow \psi$ , encoding dependencies between the action variables of other agents. Such dependencies

could be useful for modeling suspected collusion, which we will not consider in this paper. However, we do not demand that the beliefs of an agent are correct, i.e. we do not assume that each agent considers the actual game possible.

**Definition 10 (BG with Incomplete Information).** A Boolean game with incomplete information (BGI) is a tuple  $G = (\Phi_1, \dots, \Phi_n, \Gamma_1, \dots, \Gamma_n, \mathcal{K}_1, \dots, \mathcal{K}_n)$  with  $\Phi_1, \dots, \Phi_n, \Gamma_1, \dots, \Gamma_n$  as before and  $\mathcal{K}_i = \{\mathcal{K}_i^1, \dots, \mathcal{K}_i^n\}$ , where  $\mathcal{K}_i^j$  is a goal-KB of  $i$  w.r.t.  $j$ .

Let us now consider how to compute the necessity and possibility that agent  $j$  plays a best response (BR) in the strategy profile  $\nu$  according to agent  $i$ . First note that each  $\nu \in \mathcal{V}$  corresponds unambiguously to a formula  $\varphi_\nu$  in  $L_\Phi$  in the following way:

$$\varphi_\nu = \bigwedge \{p \mid p \in \nu\} \wedge \bigwedge \{\neg p \mid p \in \Phi \setminus \nu\}$$

We also introduce the following notations:

$$\begin{aligned} \varphi_{\nu_{-j}} &= \bigwedge \{p \mid p \in \nu \cap (\Phi \setminus \Phi_j)\} \wedge \bigwedge \{\neg p \mid p \in (\Phi \setminus \Phi_j) \setminus \nu\} \\ \varphi_{\nu_j} &= \bigwedge \{p \mid p \in \nu_j \cap \Phi_j\} \wedge \bigwedge \{\neg p \mid p \in \Phi_j \setminus \nu_j\} \end{aligned}$$

Note that  $\varphi_{\nu_{-j}}$  is equivalent with  $\bigvee \{\varphi_{(\nu_{-j}, \nu'_j)} \mid \nu'_j \subseteq \Phi_j\}$ .

Agent  $j$  plays a BR in the strategy profile  $\nu$  iff for every alternative strategy  $\nu'_j \subseteq \Phi_j$  it holds that  $u_j(\nu) \geq u_j(\nu_{-j}, \nu'_j)$ . Essentially this boils down to the fact that, for some  $k \in \{0, \dots, p\}$ ,  $u_j(\nu) \geq k$  and  $\forall \nu'_j \subseteq \Phi_j : u_j(\nu_{-j}, \nu'_j) \leq k$ . Note that for  $k = 0$ , the first condition is always fulfilled. Similarly, for  $k = p$ , the second condition becomes trivial. Similarly, agent  $j$  plays no BR in  $\nu$  iff there exists a  $\nu'_j \subseteq \Phi_j$  such that  $u_j(\nu) < u_j(\nu_{-j}, \nu'_j)$ . This means that, for all  $k \in \{0, \dots, p\}$ ,  $u_j(\nu) < k$  or  $\exists \nu'_j \subseteq \Phi_j : u_j(\nu_{-j}, \nu'_j) > k$ . The possibility of agent  $j$  playing a BR is dual to the necessity of agent  $j$  playing no BR. These insights motivate the following definition.

**Definition 11.** Let  $i, j \in N$  be two agents in a BGI  $G$  and let  $\nu$  be a strategy profile of  $G$ . We denote  $g_j^{p+1} = \top$  and  $g_j^0 = \perp$  for every  $j$ . We say that  $j$  plays a BR in  $\nu$  with necessity  $\lambda$  according to  $i$ , written  $BR_i^n(j, \nu) = \lambda$ , iff  $\lambda$  is the greatest value in  $[0, 1]$  for which there exists some  $k \in \{0, \dots, p\}$  such that the following two conditions are satisfied:

1.  $\mathcal{K}_i^j \vdash (\varphi_\nu \rightarrow g_j^{k+1}, \lambda)$
2.  $\mathcal{K}_i^j \vdash ((\varphi_{\nu_{-j}} \wedge \neg \varphi_{\nu_j}) \rightarrow \neg g_j^k, \lambda)$

Let  $\lambda^*$  be the smallest value greater than  $1 - \lambda$  which occurs in  $\mathcal{K}_i^j$ . Agent  $i$  believes it is possible to degree  $\lambda$  that agent  $j$  plays a BR in  $\nu$ , written  $BR_i^p(j, \nu) = \lambda$ , iff  $\lambda$  is the greatest value in  $]0, 1]$  for which there exists some  $k \in \{0, \dots, p\}$  such that the following two conditions are satisfied:

1.  $\mathcal{K}_i^j \not\vdash (\varphi_\nu \rightarrow \neg g_j^{k+1}, \lambda^*)$
2.  $\forall \nu'_j \subseteq \Phi_j : \mathcal{K}_i^j \not\vdash (\varphi_{(\nu_{-j}, \nu'_j)} \rightarrow g_j^k, \lambda^*)$

If no such  $\lambda$  exists, then  $BR_i^p(j, \nu) = 0$ .

Importantly, the syntax in Definition 11 allows to express the certainty or possibility that an agent plays a BR, from the point of view of another agent. This forms an important base from which to define interesting solution concepts or measures in BGIs. In this paper, we introduce the following measures that respectively reflect to what degree all agents believe it is necessary and possible that  $\nu$  is a PNE.

**Definition 12.** Let  $G$  be a BGI. For every strategy profile  $\nu$ , we define the measures  $PNE^n$  and  $PNE^p$  as:

$$PNE^n(\nu) = \min_{i \in N} \min_{j \in N} BR_i^n(j, \nu), \quad PNE^p(\nu) = \min_{i \in N} \min_{j \in N} BR_i^p(j, \nu)$$

If we assume that all agents know their own goal, then  $BR_i^n(i, \nu) = BR_i^p(i, \nu) = 0$  if  $\nu$  is not a PNE. Consequently, if  $\nu$  is not a PNE, then we have  $PNE^n(\nu) = PNE^p(\nu) = 0$ . Note that the measures from Definition 12 induce a total ordering on  $\mathcal{V}$ , so there always exists a  $\nu \in \mathcal{V}$  such that  $PNE^n$  or  $PNE^p$  is maximal.

*Example 4 (continued).* Let  $G$  be the BGI with the aforementioned goal-KBs and assume that Bob and Alice know their own goals. It can be computed that

	$\emptyset$	$\{a\}$	$\{b\}$	$\{a, b\}$
$\min_{j \in N} BR_1^n(j, \cdot)$	0.3	0	0	0.8
$\min_{j \in N} BR_2^n(j, \cdot)$	0.9	0	0	0.9
$PNE^n(\cdot)$	0.3	0	0	0.8

The strategy profile  $\{a, b\}$  has the highest value for  $PNE^n$ . Note that if Bob had the ‘dual’ beliefs of Alice, i.e.  $\mathcal{K}_2^1 = \{((\neg a \wedge \neg b) \rightarrow g_1^1, 0.8), ((a \wedge b) \rightarrow g_1^1, 0.3), (\neg a \rightarrow g_1^2, 1)\}$ , then  $\emptyset$  and  $\{a, b\}$  both had value 0.3 for  $PNE^n$ .

In [6], we showed that many solution concepts for BGs can be found by using a reduction to answer set programming. The concepts in this section, such as  $PNE^n$ , can be computed using a straightforward generalization of the idea in [6].

### 3.3 Soundness and Completeness

In this section, we show that the solution concepts for BGIs that were introduced in Section 3.2 indeed correspond to their semantic counterparts from Section 3.1. The classical theory  $\{\gamma_j^k \leftrightarrow g_j^k \mid k \in \{1, \dots, p\}\}$  associated with the goal base  $\Gamma_j = \{\gamma_j^1; \dots; \gamma_j^p\} \in \mathcal{G}$  is denoted as  $T_j$ . A possibility distribution  $\pi_i^j$  on  $\mathcal{G}$  can be associated with  $\mathcal{K}_i^j$  in the following natural way, inspired by [3], with  $\max \emptyset = 0$ :

$$\pi_i^j(\Gamma_j) = 1 - \max\{\alpha_l \mid (\varphi_l, \alpha_l) \in \mathcal{K}_i^j, T_j \not\models \varphi_l\} \quad (1)$$

Intuitively, the higher the certainty of the formulas violated by  $\Gamma_j$ , the lower the possibility of  $\Gamma_j$  being the real goal base of agent  $j$  according to agent  $i$ . Note that if we make the reasonable assumption that an agent knows its own goals, then  $\pi_i^i$  maps all elements of  $\mathcal{G}$  to 0 except the real goal base of  $i$ , which is mapped to 1. Given the BGI  $G$  and using the possibility distributions on  $\mathcal{G}$  for every  $j$ , we can define a possibility distribution  $\pi_i^G$  on the set of possible BGs  $\mathcal{BG}$ :

$$\pi_i^G(G') = \min_{j \in N} \pi_i^j(\Gamma_j^{G'})$$

with  $\Gamma_j^{G'}$  the goal base of agent  $j$  in the BG  $G'$ . This possibility distribution is the natural semantic counterpart of the BGI  $G$ . We now show that these possibility distributions  $\pi_i^G$  allow us to interpret the solution concepts that have been defined syntactically in Section 3.2 as instances of the semantically defined solution concepts from Section 3.1. This is formalized in the following proposition and corollary. We use the notation  $\text{br}_j(\nu, \Gamma_j)$  for the propositional variable corresponding to “agent  $j$  with goal base  $\Gamma_j$  plays a best response in  $\nu$ ”.

**Proposition 1.** *For every  $\nu \in \mathcal{V}$ ,  $i, j \in N$  and  $\lambda \in ]0, 1]$ , it holds that*

$$\text{BR}_i^n(j, \nu) \geq \lambda \Leftrightarrow \forall \Gamma_j \in \mathcal{G} : \neg \text{br}_j(\nu, \Gamma_j) \Rightarrow \pi_i^j(\Gamma_j) \leq 1 - \lambda \quad (2)$$

$$\text{BR}_i^p(j, \nu) \geq \lambda \Leftrightarrow \exists \Gamma_j \in \mathcal{G} : \text{br}_j(\nu, \Gamma_j) \wedge \pi_i^j(\Gamma_j) \geq \lambda \quad (3)$$

**Corollary 1.** *Let us denote the possibility and necessity measure associated with  $\pi_i^G$  as  $\Pi_i^G$  and  $N_i^G$ . For every  $\nu \in \mathcal{V}$  it holds that*

$$N_i^G(\{G' \in \mathcal{BG} \mid \nu \text{ is a PNE in } G'\}) = \min_{j \in N} \text{BR}_i^n(j, \nu) \quad (4)$$

$$\Pi_i^G(\{G' \in \mathcal{BG} \mid \nu \text{ is a PNE in } G'\}) = \min_{j \in N} \text{BR}_i^p(j, \nu) \quad (5)$$

Consequently, it holds that:

$$\text{nec}_G(\{G' \in \mathcal{BG} \mid \nu \text{ is a PNE in } G'\}) = \text{PNE}^n(\nu)$$

$$\text{poss}_G(\{G' \in \mathcal{BG} \mid \nu \text{ is a PNE in } G'\}) = \text{PNE}^p(\nu)$$

Before we prove Proposition 1 and Corollary 1, a lemma is stated which deals with the construction of specific goal bases in  $\mathcal{G}$ , given the knowledge about these goal bases.

**Lemma 1.** *Given a goal-KB  $\mathcal{K}_i^j$ , there exists a goal base  $\Gamma_j \in \mathcal{G}$  such that  $\pi_i^j(\Gamma_j) = 1$ .*

*Proof (Sketch).* It is easily verified that the goal base  $\Gamma_j = (\gamma_j^1; \dots; \gamma_j^p)$  with  $\gamma_j^k$  the CNF of  $\bigvee \{\varphi \mid \varphi \in L_\Phi, \exists \lambda > 0 : \mathcal{K}_i^j \vdash (\varphi \rightarrow g_j^k, \lambda)\}$  meets the condition  $\pi_i^j(\Gamma_j) = 1$ .

Note that the construction of  $\Gamma_j$  relies on the (constraint) syntax of the formulas in  $\mathcal{K}_i^j$ .

We now prove Proposition 1.

$\Rightarrow$  of (2) We prove this by contraposition. Suppose there exists a  $\Gamma_j \in \mathcal{G}$  such that  $j$  plays no BR in  $\nu$  given  $\Gamma_j$  and  $\pi_i^j(\Gamma_j) > 1 - \lambda$ . Taking (1) into account, the latter implies that  $\forall (\varphi_l, \alpha_l) \in \mathcal{K}_i^j : T_j \not\models \varphi_l \Rightarrow \alpha_l < \lambda$ . By definition 11,  $\text{BR}_i^n(j, \nu) \geq \lambda$  implies that there exists a  $k' \in \{0, \dots, p\}$  such that  $\mathcal{K}_i^j \vdash (\varphi_\nu \rightarrow g_j^{k'+1}, \lambda)$  and  $\mathcal{K}_i^j \vdash ((\varphi_{\nu-j} \wedge \neg \varphi_{\nu_j}) \rightarrow \neg g_j^{k'}, \lambda)$ . It follows that  $T_j \models \varphi_\nu \rightarrow g_j^{k'+1}$  and  $T_j \models (\varphi_{\nu-j} \wedge \neg \varphi_{\nu_j}) \rightarrow \neg g_j^{k'}$ . Consequently, by definition of  $T_j$ , if  $k' \in \{1, \dots, p-1\}$ , it holds that  $T_j \models \varphi_\nu \rightarrow \gamma_j^{k'+1}$  and  $T_j \models (\varphi_{\nu-j} \wedge \neg \varphi_{\nu_j}) \rightarrow \neg \gamma_j^{k'}$ . This means that  $j$  does play a BR in  $\nu$  since the goal  $\gamma_j^{k'+1}$  is satisfied in  $\nu$  and for every alternative strategy of  $j$ ,  $\gamma_j^{k'}$  is not satisfied. If  $k' = p$  or  $k' = 0$  then  $j$ 's utility is resp. 0 or  $p$  for every alternative strategy of  $j$ . In any case, agent  $j$  with goal base  $\Gamma_j$  plays a BR in  $\nu$ .

$\Leftarrow$  of (2) Suppose that  $\text{BR}_i^n(j, \nu) < \lambda$ , i.e. for every  $k \in \{0, \dots, p\}$  either  $\mathcal{K}_i^j \not\vdash (\varphi_\nu \rightarrow g_j^{k+1}, \lambda)$  or  $\mathcal{K}_i^j \not\vdash ((\varphi_{\nu-j} \wedge \neg \varphi_{\nu_j}) \rightarrow \neg g_j^k, \lambda)$ . Let  $k'$  be the greatest index for

which  $\mathcal{K}_i^j \not\vdash (\varphi_\nu \rightarrow g_j^{k'}, \lambda)$ . Note that  $k' \geq 1$  since  $g_j^0 = \perp$ . Construct a goal base  $\Gamma_j = (\gamma_j^1; \dots; \gamma_j^p)$  with  $\gamma_j^k$  defined as the CNF of the formula

$$\bigvee \{ \varphi \mid \varphi \in L_\Phi, \mathcal{K}_i^j \vdash (\varphi \rightarrow g_j^k, \lambda) \} \vee \left( \bigwedge \{ \varphi \mid \varphi \in L_\Phi, \mathcal{K}_i^j \vdash (\varphi \leftarrow g_j^k, \lambda) \} \wedge \neg \varphi_\nu \right)$$

for  $k \leq k'$ , and  $\gamma_j^k$  defined as the CNF of the formula

$$\bigvee \{ \varphi \mid \varphi \in L_\Phi, \mathcal{K}_i^j \vdash (\varphi \rightarrow g_j^k, \lambda) \} \vee \left( \bigwedge \{ \varphi \mid \varphi \in L_\Phi, \varphi \neq \top, \mathcal{K}_i^j \vdash (\varphi \leftarrow g_j^k, \lambda) \} \right)$$

for  $k > k'$ . One can straightforwardly check that  $\Gamma_j \in \mathcal{G}$  and  $\pi_i^j(\Gamma_j) > 1 - \lambda$  by checking that for every formula  $(\varphi, \alpha) \in \mathcal{K}_i^j$  with  $\alpha \geq \lambda$ , it holds that  $T_j \models \varphi$ . Moreover, one can verify that  $j$  does not play a BR in  $\nu$  with the constructed  $\Gamma_j$  (note that this would not be guaranteed by the goal base constructed in the proof of Lemma 1).

$\Rightarrow$  of (3) Analogous to the proof of “ $\Leftarrow$  of (2)”.

$\Leftarrow$  of (3) We prove directly that  $\text{BR}_i^p(j, \nu) \geq \lambda$ , i.e.  $\exists k \in \{0, \dots, p\}$  such that  $\mathcal{K}_i^j \not\vdash (\varphi_\nu \rightarrow \neg g_j^{k+1}, \lambda^*)$  and  $\forall \nu'_j : \mathcal{K}_i^j \not\vdash (\varphi_{(\nu_{-j}, \nu'_j)} \rightarrow g_j^k, \lambda^*)$ . By assumption, there exists a  $\Gamma_j$  such that  $j$  plays a BR in  $\nu$  and  $\pi_i^j(\Gamma_j) \geq \lambda$ . The former means that for some  $k' \in \{0, \dots, p\}$ ,  $T_j \models \varphi_\nu \rightarrow \gamma_{k'+1}^j$  and  $\forall \nu'_j : T_j \models \varphi_{(\nu_{-j}, \nu'_j)} \rightarrow \neg \gamma_{k'}^j$ . Since  $T_j \models \gamma_i^j \leftrightarrow g_j^i$ , it then holds that  $T_j \models \varphi_\nu \rightarrow g_j^{k'+1}$ . Since by definition  $\varphi_\nu \not\models \perp$ ,  $T_j \not\models \perp$  and  $T_j \not\models \neg \varphi_\nu$ , it follows that  $T_j \not\models \varphi_\nu \rightarrow \neg g_j^{k'+1}$ . The assumption that  $\pi_i^j(\Gamma_j) \geq \lambda$  implies that  $\forall (\varphi_l, \alpha_l) \in \mathcal{K}_i^j : T_j \not\models \varphi_l \Rightarrow \alpha_l \leq 1 - \lambda$ . It follows that  $\mathcal{K}_i^j \not\vdash (\varphi_\nu \rightarrow \neg g_j^{k'+1}, \lambda^*)$ . Analogously, we can prove that  $\forall \nu'_j : T_j \models \varphi_{(\nu_{-j}, \nu'_j)} \rightarrow \neg \gamma_{k'}^j$  implies that  $\forall \nu'_j : \mathcal{K}_i^j \not\vdash (\varphi_{(\nu_{-j}, \nu'_j)} \rightarrow g_j^{k'}, \lambda^*)$ .

We now prove (4) of Corollary 1. The proof of (5) is analogous and the rest of Corollary 1 follows immediately.

**Proof of (4)** By definition,  $\min_{j \in N} \text{BR}_i^n(j, \nu) \geq \lambda$  iff  $\text{BR}_i^n(j, \nu) \geq \lambda$  for every  $j \in N$ . We proved that the latter is equivalent with  $\forall \Gamma_j \in \mathcal{G} : j \text{ no BR in } \nu \Rightarrow \pi_i^j(\Gamma_j) \leq 1 - \lambda$ . We first prove that this implies that for all  $G' \in \mathcal{BG}$  it holds that  $\pi_i^G(G') \leq 1 - \lambda$  if  $\nu$  is no PNE in  $G'$ . By definition, this means that  $N_i^G(\{G' \in \mathcal{BG} \mid \nu \text{ is a PNE in } G'\}) \geq \lambda$ . Take an arbitrary  $G'$  such that  $\nu$  is no PNE in  $G'$ . Then there exists some  $j$  who plays no BR in  $\nu$  if its goal base is  $\Gamma_j^{G'}$ . By assumption, this implies  $\pi_i^j(\Gamma_j^{G'}) \leq 1 - \lambda$ , which implies  $\pi_i^G(G') \leq 1 - \lambda$  by definition. We now prove the opposite direction. Take an arbitrary  $j$  and  $\Gamma_j$  such that  $j$  plays no BR in  $\nu$  with the goal base of  $j$  equal to  $\Gamma_j$ . Using Lemma 1, we can construct a  $G' \in \mathcal{BG}$  such that  $\Gamma_j^{G'} = \Gamma_j$  and  $\pi_{j'}^i(\Gamma_{j'}^{G'}) = 1$  for every  $j' \neq j$ . Obviously  $\nu$  is no PNE in  $G'$  since  $j$  plays no BR. By assumption and definition of  $N_i^G$ , it holds that  $\pi_i^G(G') \leq 1 - \lambda$ . Since  $\pi_{j'}^i(\Gamma_{j'}^{G'}) = 1$  for every  $j' \neq j$ , it follows that  $\pi_i^j(\Gamma_j) \leq 1 - \lambda$ . Due to Proposition 1, we proved that  $\text{BR}_i^n(j, \nu) \geq \lambda$ . Since  $j$  is arbitrary, it follows that  $\min_{j \in N} \text{BR}_i^n(j, \nu) \geq \lambda$ .

*Example 5.* Recall the scenario of Example 2. We define the BGI  $G$ . Since Bob has perfect knowledge of Alice’s preferences, his goal-KB can be modeled as  $\mathcal{K}_2^1 = \mathcal{K}_1^1 = \{((a \leftrightarrow b) \leftrightarrow g_1^1, 1), (((a \leftrightarrow b) \vee \neg a) \leftrightarrow g_1^2, 1)\}$ . Alice is certain that Bob wants to attend a sports game together, or attend the game on his own if attending it together is not possible. This can be captured by the goal-KB  $\mathcal{K}_1^2 = \{((a \wedge b) \leftrightarrow g_2^1, 1), (b \leftrightarrow g_2^2, 1)\}$ . It is easy to see that  $\pi_1^G$  and  $\pi_2^G$  correspond to the possibility distributions

$\pi_1$  and  $\pi_2$  described in Example 2. Despite Alice's incorrect beliefs, Bob and Alice are both certain that attending a sports game together is a PNE, since  $nec_G(\{G' \in \mathcal{BG} \mid \{a, b\} \text{ is a PNE in } G'\}) = PNE^n(\{a, b\}) = 1$ . Contrary to Alice, Bob knows that going to the theatre together is a PNE as well.

An interesting question is how the agents' beliefs can influence the proposals they can make in e.g. bargaining protocols. Suppose for instance that Alice wants to make Bob a suggestion, then based on her beliefs, it would be rational to suggest to attend a sports game together. Bob would then rationally agree, based on his beliefs. However, if Bob were to make a proposal, he can choose between two rational suggestions: attending a sports game together or going to the theatre together. If he would do the latter, Alice would know that her beliefs are incorrect, assuming Bob behaves rationally. In future research, we will investigate these strategical interactions and how they allow agents to revise their beliefs. Other research possibilities lie in manipulating BGIs through communication, for instance through announcements, as investigated for BGs with environment variables [11]. Another option is to extend the BGI framework, allowing agents to also reason about the beliefs of other agents, although this is likely to lead to an increase in computational complexity.

## 4 Decision problems

The decision problems associated with BGIs and the  $PNE^x$  measures are investigated.

**Proposition 2.** *Let  $G$  be a BGI and  $\lambda \in ]0, 1]$ . The following decision problems are  $\Sigma_2^P$ -complete:*

1. *Does there exist a strategy profile  $\nu$  with  $PNE^n(\nu) \geq \lambda$ ?*
2. *Does there exist a strategy profile  $\nu$  with  $PNE^p(\nu) \geq \lambda$ ?*

*Proof.* Hardness of 1 and 2 Both problems are  $\Sigma_2^P$ -hard since they contain the  $\Sigma_2^P$ -complete problem to decide whether a BG has a PNE as a special case. Indeed, when  $G$  is a BG, we can construct a BGI in which all agents have complete knowledge of each others goals. Then  $PNE^n(\nu)$  and  $PNE^p(\nu)$  coincide and take values in  $\{0, 1\}$ , depending on whether  $\nu$  is a PNE or not. Consequently,  $G$  has a PNE iff there exists a  $\nu$  with  $PNE^n(\nu) = PNE^p(\nu) \geq \lambda$ .

Completeness of 1 We can decide the problem by first guessing a strategy profile  $\nu$ . Checking whether  $PNE^n(\nu) \geq \lambda$  means checking whether  $BR_i^n(j, \nu) \geq \lambda$  for every  $i, j \in N$ . The latter involves checking possibilistic entailments, which can be done in constant time using an NP-oracle. Therefore, the decision problem is  $\Sigma_2^P$ -complete.

Completeness of 2 We can decide the problem by first guessing a strategy profile  $\nu$ . Checking whether  $PNE^p(\nu) \geq \lambda$  means checking whether  $BR_i^p(j, \nu) \geq \lambda$  for every  $i, j \in N$ . To see that the latter can be reduced to checking a polynomial number of possibilistic entailments, we need to rewrite the condition that  $\forall \nu'_j \subseteq \Phi_j : \mathcal{K}_i^j \not\vdash (\varphi(\nu_{-j}, \nu'_j) \rightarrow g_j^k, \lambda^*)$ . To this end, we define  $\mathcal{K}^k$ , for every  $k \in \{1, \dots, p\}$ , as the KB  $\mathcal{K}_i^j$  in which all formulas defining necessary and/or sufficient conditions for  $g_j^k$  are

preserved; all formulas with necessary conditions for  $g_j^l$  ( $l \geq k$ ) are translated into necessary conditions for  $g_j^k$  by replacing  $(\varphi \rightarrow g_j^l, \alpha)$  by  $(\varphi \rightarrow g_j^k, \alpha)$ ; all formulas with sufficient conditions for  $g_j^l$  ( $l \leq k$ ) are translated into sufficient conditions for  $g_j^k$  by replacing  $(\varphi \leftarrow g_j^l, \alpha)$  by  $(\varphi \leftarrow g_j^k, \alpha)$ ; all other formulas are deleted. Then it holds

$$\begin{aligned}
& \forall \nu'_j \subseteq \Phi_j : \mathcal{K}_i^j \not\models (\varphi_{(\nu_{-j}, \nu'_j)} \rightarrow g_j^k, \lambda^*) \\
& \Leftrightarrow \forall \nu'_j \subseteq \Phi_j : \mathcal{K}^k \not\models (\varphi_{(\nu_{-j}, \nu'_j)} \rightarrow g_j^k, \lambda^*) \\
& \Leftrightarrow \forall \nu'_j \subseteq \Phi_j : \mathcal{K}_{1-\lambda}^k \not\models \varphi_{(\nu_{-j}, \nu'_j)} \rightarrow g_j^k \\
& \Leftrightarrow \forall \nu'_j \subseteq \Phi_j : \mathcal{K}_{1-\lambda}^k \text{ and } \varphi_{(\nu_{-j}, \nu'_j)} \text{ and } \neg g_j^k \text{ are consistent} \\
& \Leftrightarrow \forall \nu'_j \subseteq \Phi_j : \mathcal{K}_{1-\lambda}^k \text{ and } \varphi_{\nu'_j} \text{ are consistent}
\end{aligned}$$

where  $\mathcal{K}_{1-\lambda}^k$  is obtained from  $\mathcal{K}_{1-\lambda}^k$  by replacing each occurrence of  $g^k$  by  $\perp$  and each occurrence of  $p \in \Phi \setminus \Phi_j$  by its truth value ( $\top$  or  $\perp$ ) in  $\nu$ . The last condition is equivalent with  $\mathcal{K}_{1-\lambda}^k$  being a tautology, which can be checked with a SAT-solver, i.e. in constant time with an NP-oracle.

The result of Proposition 2 shows that the complexity for the introduced measures does not increase compared to PNEs of BGs, since deciding whether a BG has a PNE is also  $\Sigma_2^P$ -complete. Moreover, given the experimental results reported in [6] for standard BGs, it seems plausible that a reduction to answer set programming would support an efficient computation of solutions for medium sized games.

## 5 Conclusion

We introduced the first BG framework that allows agents to be uncertain about the other agents' goals. We have argued that such a scenario can naturally be modeled by associating with each agent a possibility distribution over the universe of all possible games (given the considered partition of action variables). While this allows us to define a variety of solution concepts in a natural way, this semantic approach is not useful in practice, due to the exponential size of these possibility distributions. Therefore, we also proposed a syntactic counterpart, which avoids exponential representations by relying on standard possibilistic logic inference, and can be implemented by reduction to answer set programming. Our main result is that this syntactic characterization indeed corresponds to the intended semantic definitions. We furthermore showed that the computational complexity of reasoning with our Boolean games with incomplete information remains at the second level of the polynomial hierarchy. The present framework leads to several interesting avenues for future work. First, the approach could be generalized for taking into account prior beliefs about the likely behaviour of other players (e.g. for modeling collusion) and/or for modeling situations where agents may be uncertain about the actions that are being played by other agents. Moreover, it seems of interest to analyse the effect of adding communication to the framework, by allowing agents to strategically ask questions or make proposals to each other in order to reduce uncertainty or as part of a bargaining process.

## References

1. Ågotnes, T., Harrenstein, P., van der Hoek, W., Wooldridge, M.: Boolean games with epistemic goals. In: Logic, Rationality, and Interaction, LNCS, vol. 8196, pp. 1–14. Springer (2013)
2. Ågotnes, T., Harrenstein, P., van der Hoek, W., Wooldridge, M.: Verifiable equilibria in Boolean games. In: Proc. IJCAI '13 (2013)
3. Benferhat, S., Kaci, S.: Logical representation and fusion of prioritized information based on guaranteed possibility measures: Application to the distance-based merging of classical bases. *Artificial Intelligence* 148, 291–333 (2003)
4. Bonzon, E., Lagasquie-Schiex, M.C., Lang, J.: Compact preference representation for Boolean games. In: Proc. PRICAI '06, pp. 41–50. Springer-Verlag (2006)
5. Bonzon, E., Lagasquie-Schiex, M.C., Lang, J., Zanuttini, B.: Boolean games revisited. In: Proc. ECAI '06, pp. 265–269. ACM (2006)
6. De Clercq, S., Bauters, K., Schockaert, S., De Cock, M., Nowé, A.: Using answer set programming for solving Boolean games. In: Proc. KR '14, to appear (2014)
7. Dubois, D., Lang, J., Prade, H.: Possibilistic logic. In: *Handbook of Logic for Artificial Intelligence and Logic Programming*, vol. 3, pp. 439–513. Oxford University Press (1994)
8. Dubois, D., Prade, H.: Epistemic entrenchment and possibilistic logic. *Artificial Intelligence* 50(2), 223–239 (1991)
9. Dubois, D., Prade, H.: Possibilistic logic: a retrospective and prospective view. *Fuzzy Sets and Systems* 144, 3–23 (2004)
10. Dunne, P., van der Hoek, W., Kraus, S., Wooldridge, M.: Cooperative Boolean games. In: Proc. AAMAS '08, vol. 2, pp. 1015–1022. IFAAMAS (2008)
11. Grant, J., Kraus, S., Wooldridge, M., Zuckerman, I.: Manipulating Boolean games through communication. In: Proc. IJCAI '11 (2011)
12. Harrenstein, P., van der Hoek, W., Meyer, J.J., Witteveen, C.: Boolean games. In: Proc. TARK '01, pp. 287–298. MKP Inc. (2001)
13. Marchioni, E., Wooldridge, M.: Łukasiewicz games. In: Proc. AAMAS '14, pp. 837–844 (2014)
14. Osborne, M., Rubinstein, A.: *A Course in Game Theory*. MIT press (1994)
15. Schelling, T.: *The strategy of conflict*. Oxford University Press (1960)
16. Schockaert, S., Prade, H.: Solving conflicts in information merging by a flexible interpretation of atomic propositions. *Artificial Intelligence* 175(11), 1815–1855 (2011)