

Multilateral Negotiation in Boolean Games with Incomplete Information using Generalized Possibilistic Logic*

Sofie De Clercq

Dept. of Applied Maths, CS and Statistics
Ghent University, Ghent, Belgium
{sofier.declercq, martine.decock}@ugent.be

Steven Schockaert

School of Computer Science and Informatics
Cardiff University, Cardiff, UK
schockaerts1@cardiff.ac.uk

Ann Nowé

Computational Modeling Lab
Vrije Universiteit Brussel, Brussels, Belgium
ann.nowe@vub.ac.be

Martine De Cock

Center for Data Science
University of Washington, Tacoma, USA
mdecock@uw.edu

Abstract

Boolean games are a game-theoretic framework in which propositional logic is used to describe agents' goals. In this paper we investigate how agents in Boolean games can reach an efficient and fair outcome through a simple negotiation protocol. We are particularly interested in settings where agents only have incomplete knowledge about the preferences of others. After explaining how generalized possibilistic logic can be used to compactly encode such knowledge, we analyze how a lack of knowledge affects the agreement outcome. In particular, we show how knowledgeable agents can obtain a more desirable outcome than others.

1 Introduction and Related Work

Boolean games (BGs) are a strategic framework which uses propositional logic to formalize the goals of agents [Harrenstein *et al.*, 2001]. Each agent controls the truth value of a subset of the atoms which build up these goals. In this paper, we investigate BGs in which the agents are uncertain about the preferences of the other agents. For example, suppose Bob and Alice are planning their Sunday afternoon: they can go to the beach or the forest or stay at home, and Alice can bring the dog or leave it at home. Since being a couple does not imply having identical preferences nor knowing exactly each other's preferences, Bob and Alice will have to compromise under incomplete information. However, they are not completely ignorant about each other's goals; for instance, Bob knows that Alice loves the dog and Alice knows that Bob loves the beach. In this paper, we investigate how they can use such knowledge to reach an agreement through negotiation. Moreover, we explore the link between having information about other agents' goals and obtaining a satisfactory agreement. To the best of our knowledge, our process is the

first multilateral negotiation protocol for BGs that takes uncertainty w.r.t. the other agents' goals into account.

Although uncertainty in games has been studied extensively (see e.g. [Osborne and Rubinstein, 1994]), the literature on BGs with incomplete information is currently limited. In one framework, agents have beliefs about the values of environment variables, which are not controlled by any of the agents [Grant *et al.*, 2011]. Another approach limits the set of action variables of which an agent can observe the assigned values [Ågotnes *et al.*, 2013b]. The latter framework has been extended to epistemic BGs [Ågotnes *et al.*, 2013a], describing goals by means of a multi-agent epistemic modal logic. However, in the latter approach agents are still fully aware of each others' goals, i.e. Ågotnes *et al.* [2013a] consider agents whose goal is to obtain a particular epistemic state. To the best of our knowledge, the only work on BGs with uncertainty w.r.t. the agents' goals is [De Clercq *et al.*, 2014], which uses possibilistic knowledge bases to formalize the beliefs of the agents about the goals of others.

In this paper, we propose the use of generalized possibilistic logic (GPL) [Dubois *et al.*, 2014] to define BGs with incomplete information. Note that this is different from the use of uncertainty models, such as the Bayesian approaches commonly applied to study games with uncertainty (see e.g. [Harsanyi, 1967]), which can be used e.g. to model beliefs about the characteristics of other agents. Although De Clercq *et al.* [2014] also use possibilistic logic to model uncertainty in BGs, our motivation differs: whereas the former work uses possibilistic logic to encode graded beliefs about other agents' goals, we use possibilistic logic to compactly describe agents' preferences, and generalized possibilistic logic to describe incomplete knowledge about another agent's preferences. Our use of GPL also differs from approaches such as CP-nets, which also aim to compactly model preferences, but only capture a single preference structure and are thus less suitable for modeling incomplete information. To the best of our knowledge, this is the first research on using GPL to model uncertainty about other agents' preferences.

We then investigate how knowledge about other agents' goals can be used for multilateral negotiation. Negotiating

*This research is funded by the Research Foundation Flanders (FWO).

allows agents in a strategic setting to settle on an agreement outcome, where agents can be human as well as artificial, such as computers, robots or self-driving cars. A multilateral bargaining protocol in BGs with complete information has been investigated in [Dunne *et al.*, 2008], showing that, when the logical structure of the goals is restricted, the protocol is guaranteed to end in a Pareto optimal outcome, i.e. no agent can improve its position without another agent being worse off. In this paper, we propose a protocol which converges to an acceptable agreement without restrictions on the game structure. Moreover, under complete information our protocol always results in a discrim optimal outcome. Discrimin optimality refines Pareto optimality [Dubois *et al.*, 1997], and while the latter indeed ensures *efficiency*, it is often not sufficient to characterize desirable outcomes [Bouveret and Lang, 2008]. Suppose, for instance, that two agents are negotiating in a situation with two Pareto optimal outcomes, with utility vectors (1, 0.2) respectively (0.5, 0.6), where utility reflects the degree of satisfaction of the agents. A natural concept arising in negotiation is *fairness*: intuitively, the latter utility vector is more fair than the former. In the literature, several notions of fairness apart from discrim optimality have been introduced and studied; we refer the interested reader to [Tun- godden, 2000] for an overview and discussion.

We investigate negotiation in BGs with incomplete information since, in strategic settings, it is often unrealistic to assume that an agent has access to another agent’s goals. Indeed, agents might deliberately conceal such information, or might not have exchanged it. The literature on bargaining is extensive, and covers a wide range of possible settings, such as discrete versus continuous bargaining (e.g. prices [Fatima *et al.*, 2002]), bilateral [Amgoud and Prade, 2003; Luo *et al.*, 2003] versus multilateral bargaining, transferable and non-transferable utility, limited versus unlimited number of responses, modeling incomplete knowledge through probability theory [Fatima *et al.*, 2002], through Cartesian products of complete knowledge problems [Bossert and Peters, 2001], or by means of possibilistic logic [Amgoud and Prade, 2003]. To the best of our knowledge, our work is the first research on negotiation that considers incomplete knowledge about other agents’ goals in a BG setting.

The paper is structured as follows. First, we give some background on BGs. In Section 3, we propose a negotiation protocol for BGs with *complete* information and characterize the agreement outcomes. Next, we explain how GPL can be used to represent knowledge about other agents’ preferences in Section 4. Then we generalize the negotiation protocol of Section 3 to BGs with *incomplete* information. We characterize the agreement outcomes, linking back to those under complete information. Additionally, we show how knowledge is crucial for an agent to reach a satisfying agreement. To conclude, we discuss several future work directions.

2 Background on Boolean Games

The logical language L_Φ associated with a finite set of atoms Φ contains the following formulas: (i) every atom of Φ , (ii) the logical constants \perp and \top , and (iii) the formulas $\neg\varphi$ and $\varphi \wedge \psi$ for every $\varphi, \psi \in L_\Phi$. As usual, we use the abbreviations

$\varphi \rightarrow \psi \equiv \neg(\varphi \wedge \neg\psi)$ and $\varphi \vee \psi \equiv \neg(\neg\varphi \wedge \neg\psi)$. An interpretation of Φ is defined as a subset ν of Φ , with the convention that all atoms in ν are interpreted as true (\top) and all atoms in $\Phi \setminus \nu$ are interpreted as false (\perp). An interpretation can be extended to L_Φ in the usual way. If a formula $\varphi \in L_\Phi$ is true in an interpretation ν , we denote this as $\nu \models \varphi$. We denote the set of interpretations in which φ is true as $\llbracket \varphi \rrbracket$.

Originally, the utilities in BGs were binary, but several extensions have been introduced to allow more general preferences. Examples are the addition of costs [Dunne *et al.*, 2008], the use of a prioritized goal base instead of a single goal [Bonzon *et al.*, 2006; De Clercq *et al.*, 2014] or the use of many-valued Łukasiewicz logic to formalize the idea of weighted goal satisfaction [Marchioni and Wooldridge, 2014]. In our paper, we use the definition of a BG as stated in [De Clercq *et al.*, 2014]. The latter is a particular case of generalized BGs [Bonzon *et al.*, 2006] in which the preference relations are total. Additionally, we incorporate a constraint δ , restricting the possible joint actions of the agents. This is a generalization of the constraints in [Bonzon *et al.*, 2012], which only restrict the individual actions of the agents.

Definition 1 (Boolean Game)

A *Boolean game (BG)* is a tuple $G = (\delta, \Phi_1, \dots, \Phi_n, \Gamma_1, \dots, \Gamma_n)$. For every agent i in $N = \{1, \dots, n\}$, Φ_i is a finite set of atoms such that $\Phi_i \cap \Phi_j = \emptyset, \forall j \neq i$. We write $\Phi = \bigcup_{i \in N} \Phi_i$. For every $i \in N$, $\Gamma_i = \{\gamma_i^1; \dots; \gamma_i^p\}$ is i ’s prioritized goal base. The formula $\gamma_i^m \in L_\Phi$ is agent i ’s goal of priority m . We assume that every agent has p priority levels and that $\delta \wedge \gamma_i^m \neq \perp$ for every $i \in N$ and $m \in \{1, \dots, p\}$. Finally, δ is a consistent formula in L_Φ , which encodes the integrity constraints of the game G .

The set Φ contains all action variables. Agent i controls Φ_i and can set these atoms to true or false. By convention, goals are ordered from high (level 1) to low priority (level p).

Example 1

Alice and Bob, who share a car, are planning their afternoon. Alice controls $\Phi_1 = \{b_A, f_A, d_A\}$ and Bob controls $\Phi_2 = \{b_B, f_B\}$. Agent i can drive to the beach (set b_i to true) or to the forest (set f_i to true). If Alice sets d_A to true, she takes the dog. The game is constrained by $\delta = \neg(b_B \wedge f_B) \wedge \neg(b_A \wedge f_A) \wedge (b_B \rightarrow \neg f_A) \wedge (b_A \rightarrow \neg f_B)$. Alice and Bob’s goal bases are

$$\begin{aligned} \Gamma_1 &= \{f_A \wedge f_B \wedge d_A; b_A \wedge b_B \wedge d_A; d_A\} \\ \Gamma_2 &= \{b_B \wedge b_A \wedge \neg d_A; b_B \wedge b_A; f_B \wedge f_A\} \end{aligned}$$

Thus Alice prefers staying at home with her dog over leaving without it. Bob prefers to take Alice to the beach without the dog. However, he still prefers to go to the beach with Alice and the dog over going to the forest with Alice, and he prefers going to the forest with Alice over all remaining possibilities.

Definition 2 (Outcome)

An interpretation of Φ is called an *outcome* of G . We denote the set of all outcomes as \mathcal{V} .

For the ease of presentation, we define a utility function that is scaled to the unit interval.

Definition 3 (Utility Function)

For each $i \in N$ and $\nu \in \mathcal{V}$, the utility of i in ν is defined as

$$u_i(\nu) = \frac{p + 1 - \min\{k \mid 1 \leq k \leq p, \nu \models \gamma_i^k \wedge \delta\}}{p}$$

with $\min \emptyset = p + 1$.

The utility takes values in $\Lambda_p = \{0, \frac{1}{p}, \frac{2}{p}, \dots, 1\}$. We will denote the vector of utility functions $(u_1(\nu), \dots, u_n(\nu))$ as $\mathcal{U}(\nu)$. In Example 1, we have for instance $\mathcal{U}(\{f_A, f_B, d_A\}) = (1, 0.33)$, $\mathcal{U}(\{b_B, b_A\}) = (0, 1)$, $\mathcal{U}(\{b_A, b_B, d_A\}) = (0.67, 0.67)$ and $\mathcal{U}(\{f_A, b_B\}) = (0, 0)$. In the context of bargaining, it is natural that agents try to achieve an outcome that is, among others, efficient. A well-known efficiency concept is Pareto optimality.

Definition 4 (Pareto Efficiency)

For every $\nu, \nu' \in \mathcal{V}$ it holds that $\nu >_p \nu'$ iff

$$(\forall i \in N : u_i(\nu) \geq u_i(\nu')) \wedge (\exists i \in N : u_i(\nu) > u_i(\nu'))$$

We denote the set of Pareto optimal outcomes in \mathcal{V} as

$$Opt_{par} = \{\nu \in \llbracket \delta \rrbracket \mid \neg(\exists \nu' \in \mathcal{V} : \nu' >_p \nu)\}$$

Intuitively, an outcome is Pareto optimal if no agent can be better off without another agent being worse off. It is easy to see that every BG has at least one Pareto optimal outcome.

A well-known refinement of the Pareto ordering incorporating a notion of fairness is the discrimin ordering [Dubois et al., 1997]. To define it, we denote the set of agents whose utility is the same in ν and ν' as $eq(\mathcal{U}(\nu), \mathcal{U}(\nu'))$, i.e. $eq(\mathcal{U}(\nu), \mathcal{U}(\nu')) = \{i \in N \mid u_i(\nu) = u_i(\nu')\}$.

Definition 5 (Discrimin Ordering)

For every $\nu, \nu' \in \mathcal{V}$ it holds that $\nu >_d \nu'$ iff

$$\min_{j \notin eq(\mathcal{U}(\nu), \mathcal{U}(\nu'))} u_j(\nu) > \min_{j \notin eq(\mathcal{U}(\nu'), \mathcal{U}(\nu))} u_j(\nu')$$

We define the set of discrimin optimal outcomes as

$$Opt_{discr} = \{\nu \in \llbracket \delta \rrbracket \mid \neg(\exists \nu' \in \mathcal{V} : \nu' >_d \nu)\}$$

It is easy to see that $>_d$ is a strict order relation on \mathcal{V} . Straightforwardly, it holds that $Opt_{discr} \subseteq Opt_{par}$ and $Opt_{discr} \neq \emptyset$. In Example 1, (b_A, b_B, d_A) is the unique discrimin optimal outcome, although it is not the only Pareto optimal outcome.

3 Negotiating under Complete Information

We are interested in a negotiation protocol that is guaranteed to converge within a finite number of steps. Therefore, we want agents to make offers according to a negotiation rule, which assures that every offered outcome is an improvement compared to the previous one. For instance, an agent might only be allowed to make a counteroffer if no agent is worse off than in the previous offer. Obviously, this rule will lead to Pareto optimal outcomes. However, the rule is so strict that the result can hardly be called fair: the first agent simply offers the outcome which yields its personal highest utility and no other agent is allowed to make a counteroffer which lowers the first agent's utility. Suppose, for instance, that there are two possible utility vectors: $(1, 0)$ and $(0.5, 0.5)$. If the

first agent opens the negotiation with $(1, 0)$, the other agent would not be allowed to counter this offer with $(0.5, 0.5)$. To develop a fairer rule, we consider two properties that characterize a valid counteroffer. First of all, an agent is only interested in making a counteroffer if its own utility improves compared to the original offer. Second, the agents apply the *silver rule* or *ethic of reciprocity*, proposed by the Confucian Way of Humanity [Hertzler, 1934]:

One should not treat others in ways that one would not like to be treated.

In our negotiation protocol, an agent reasons as follows: if I do not accept an offer of utility k , I should not lower another agent's utility to k or less in order to improve my own. Therefore, if an agent decides to lower other agents' utilities, it should offer more than k . We formally define the set $co(i, \nu)$ of agent i 's legal counteroffers to ν as follows ($i \in N, \nu \in \mathcal{V}$):

$$co(i, \nu) = \{\nu' \in \mathcal{V} \mid u_i(\nu') > u_i(\nu) \wedge \forall j \in N : u_j(\nu') < u_j(\nu) \Rightarrow u_j(\nu') > u_i(\nu)\} \quad (1)$$

We suggest the following negotiation protocol. In a given order, agents make proposals one by one. Without loss of generality, we assume that this order is $\{1, 2, \dots, n\}$.

Algorithm 1 Negotiation Protocol for BG

```

 $\nu \leftarrow \nu'$  with  $\nu' \in \llbracket \delta \rrbracket$  % Agent 1 proposes  $\nu'$ 
accepted  $\leftarrow 1$ ;  $i \leftarrow 2$ 
while accepted  $< n$  do
  if  $co(i, \nu) == \emptyset$  then
    % Agent  $i$  accepts the offer
    accepted  $\leftarrow$  accepted + 1
  else
    % Agent  $i$  rejects the offer and makes a counteroffer
     $\nu \leftarrow \nu'$  with  $\nu' \in co(i, \nu)$ 
    accepted  $\leftarrow 1$ 
  end if
   $i \leftarrow (i == n ? 1 : i + 1)$ 
end while

```

Algorithm 1 depends on a selection function to choose which $\nu' \in \llbracket \delta \rrbracket$ and $\nu' \in co(i, \nu)$ are made as an offer. However, the results discussed in the paper hold regardless of this selection. The negotiation protocol ends if an offer ν is made such that no counteroffers can be made, i.e. $\forall i \in N : co(i, \nu) = \emptyset$. We can prove that, whenever an offer is rejected, the new offer is fairer according to the discrimin ordering.

Proposition 1

For $\nu \in \mathcal{V}$, $i \in N$ and $\nu' \in co(i, \nu)$, it holds that $\nu' >_d \nu$.

Sketch. This follows immediately from the definitions of $co(i, \nu)$ and $>_d$, and the fact that either $u_j(\nu') > u_j(\nu)$ or $u_j(\nu') < u_j(\nu)$ for every $j \notin eq(\mathcal{U}(\nu), \mathcal{U}(\nu'))$. \square

It follows that every discrimin optimal outcome is accepted.

Corollary 2

For $\nu \in Opt_{discr}$ and $i \in N$ it holds that $co(i, \nu) = \emptyset$.

Conversely, we can also show that only discrimin optimal outcomes will be overall accepted.

Proposition 3

For $\nu \in \mathcal{V} \setminus Opt_{discr}$ there is an $i \in N$ with $co(i, \nu) \neq \emptyset$.

Sketch. By definition of Opt_{discr} , there exists a $\nu' \in \mathcal{V}$ such that $\nu' >_d \nu$. It is straightforward to prove that $\nu' \in co(i, \nu)$ for every agent $i \notin eq(\mathcal{U}(\nu), \mathcal{U}(\nu'))$ such that $u_i(\nu) = \min_{j \notin eq(\mathcal{U}(\nu), \mathcal{U}(\nu'))} u_j(\nu)$. \square

Note that since there are only a finite number of offers that can be made, and because each offer must strictly improve the previous offer in terms of the discrimin ordering, we know that the negotiation protocol always ends. From Corollary 2 and Proposition 3 we moreover know that the possible outcomes at the end of the negotiation protocol are exactly the discrimin optimal outcomes. This result implies that the first offering agent still has a strong advantage, as this agent can select the discrimin optimal outcome that yields the highest personal utility, which no agent is allowed to reject. For instance, if the only discrimin optimal outcomes have utility vectors (1,0.5) and (0.5,1), agent 1 should propose the former and agent 2 has no choice but to accept. If the first agent follows this strategy, the negotiation ends within one step.

Remark 1

In our protocol, it is irrelevant which agent controls which atoms. The dependence of actions implied by the constraint δ forces agents to negotiate about what actions they will undertake. In Example 1, Alice and Bob cannot decide individually to go out. However, Alice can decide to stay with the dog without violating δ . Moreover, both Alice and Bob can decide to stay home without restricting the other agent's options w.r.t. the constraint. Thus Alice is able to reach a utility of 0.33 without negotiating, and Bob will be stuck with a utility of 0. The utility vector (0.33, 0) can be viewed as the *disagreement point* [Binmore *et al.*, 1986], i.e. the utility the agents would receive if they fail to reach an agreement. This information could be added to the framework: Alice rejects everything with a lower utility than 0.33, ergo Bob should not make such offers during the negotiation. Note that we can incorporate this info in the constraint δ , demanding that the utility of every agent is greater than its *disagreement utility*.

4 Modeling Knowledge about Preferences using Generalized Possibilistic Logic

Possibilistic logic extends classical logic by associating weights with formulas. These weights were originally interpreted as degrees of certainty [Dubois *et al.*, 1994], but can also be viewed as degrees of satisfaction [Dubois *et al.*, 1996; Benferhat *et al.*, 2001], allowing the use of possibilistic logic for modeling preferences. Generalized possibilistic logic (GPL) is a recent extension of possibilistic logic, which was introduced to model incomplete knowledge about the beliefs of another agent [Dubois *et al.*, 2014]. In this paper, we use GPL to compactly encode what each agent knows about the preferences of another agent.

We define $\Lambda_p^+ = \{\frac{1}{p}, \frac{2}{p}, \dots, 1\}$. The GPL language $L_{\mathbb{N}}^p$ with $p + 1$ satisfaction levels is defined as follows:

- If $\alpha \in L_{\Phi}$ and $\lambda \in \Lambda_p^+$, then $\mathbf{N}_{\lambda}(\alpha) \in L_{\mathbb{N}}^p$.
- If $\gamma_1 \in L_{\mathbb{N}}^p$ and $\gamma_2 \in L_{\mathbb{N}}^p$, then $\neg\gamma_1$ and $\gamma_1 \wedge \gamma_2$ in $L_{\mathbb{N}}^p$.

We define the involutive function $inv : \Lambda_p^+ \rightarrow \Lambda_p^+$ with $inv(\lambda) = \frac{p+1}{p} - \lambda$. GPL uses the following abbreviations:

$$\mathbf{\Pi}_{\lambda}(\alpha) \equiv \neg \mathbf{N}_{inv(\lambda)}(\neg\alpha), \quad \mathbf{\Delta}_{\lambda}(\alpha) \equiv \bigwedge_{\nu \in \llbracket \alpha \rrbracket} \mathbf{\Pi}_{\lambda}(\varphi_{\nu})$$

where φ_{ν} for $\nu \in \mathcal{V}$ is defined as the conjunction of literals that are true in ν , i.e. $\varphi_{\nu} = \bigwedge_{\nu \models p} p \wedge \bigwedge_{\nu \not\models p} \neg p$. The semantics of GPL formulas are defined through normalized possibility distributions. In this context, a possibility distribution π is a $\mathcal{V} \rightarrow \Lambda_p$ mapping, encoding the degree to which each outcome is desirable to a given agent. A GPL knowledge base \mathcal{K} is a finite set of formulas in $L_{\mathbb{N}}^p$. It holds that π is a model of a GPL knowledge base \mathcal{K} iff π is a model of each formula in \mathcal{K} . Moreover:

- π is a model of $\mathbf{N}_{\lambda}(\alpha)$ iff $N(\alpha) \geq \lambda$;
- π is a model of $\gamma_1 \wedge \gamma_2$ iff π is a model of γ_1 and π is a model of γ_2 ;
- π is a model of $\neg\gamma_1$ iff π is not a model of γ_1 ;

where N is the necessity measure induced by π , i.e. $N(\alpha) = \min_{\nu \not\models \alpha} (1 - \pi(\nu))$. The set of all models of \mathcal{K} is denoted as $Mod(\mathcal{K})$. In this paper, we use the following established links between syntax and semantics [Dubois *et al.*, 2014]:

$$\begin{aligned} \mathcal{K} \models \mathbf{N}_{\lambda}(\alpha) &\equiv (\forall \pi \in Mod(\mathcal{K}), \forall \nu \in \mathcal{V} : (\nu \not\models \alpha) \\ &\Rightarrow \pi(\nu) \leq 1 - \lambda) \\ \mathcal{K} \models \mathbf{\Pi}_{\lambda}(\alpha) &\equiv (\exists \pi \in Mod(\mathcal{K}), \exists \nu \in \mathcal{V} : (\nu \models \alpha) \\ &\wedge \pi(\nu) \geq \lambda) \\ \mathcal{K} \models \mathbf{\Delta}_{\lambda}(\alpha) &\equiv (\forall \pi \in Mod(\mathcal{K}), \forall \nu \in \mathcal{V} : (\nu \models \alpha) \\ &\Rightarrow \pi(\nu) \geq \lambda) \end{aligned}$$

In particular, it holds for every $\nu \in \mathcal{V}$ that:

$$\begin{aligned} \mathcal{K} \models \mathbf{N}_{\lambda}(\neg\varphi_{\nu}) &\equiv \forall \pi \in Mod(\mathcal{K}) : \pi(\nu) \leq 1 - \lambda \\ \mathcal{K} \models \mathbf{\Pi}_{\lambda}(\varphi_{\nu}) &\equiv \exists \pi \in Mod(\mathcal{K}) : \pi(\nu) \geq \lambda \\ \mathcal{K} \models \mathbf{\Delta}_{\lambda}(\varphi_{\nu}) &\equiv \forall \pi \in Mod(\mathcal{K}) : \pi(\nu) \geq \lambda \end{aligned}$$

We explain the intuition of GPL formulas below, after introducing the BG framework with incomplete preference-information by means of GPL knowledge bases.

Definition 6 (BG with Incomplete Preference-Information)
A Boolean game with incomplete preference-information is a tuple $G = (\delta, \Phi_1, \dots, \Phi_n, \Gamma_1, \dots, \Gamma_n, \mathcal{K}_1, \dots, \mathcal{K}_n)$ 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 GPL knowledge base such that $\mathbf{N}_1(\delta) \in \mathcal{K}_i^j$, $Mod(\mathcal{K}_i^j) = \{u_i\}$ and $u_j \in Mod(\mathcal{K}_i^j)$ for every $i, j \in N$.

The GPL base \mathcal{K}_i^j encodes what agent i knows about the preferences of j . The assumption $\mathbf{N}_1(\delta) \in \mathcal{K}_i^j$ means that each agent knows that all agents are aware of the integrity constraint, i.e. agent i knows that the utility of agent j is 0 if the outcome violates δ . The assumption $Mod(\mathcal{K}_i^j) = \{u_i\}$

means that agent i knows its own utility. This can be accomplished if \mathcal{K}_i^i contains the formulas $\mathbf{N}_{\frac{m}{p}}(\bigvee_{l=1}^m \gamma_l^i)$ and $\Delta_{inv(\frac{m}{p})}(\gamma_i^m \wedge \delta)$ for every $m \in \{1, \dots, p\}$. Finally, the assumption $u_j \in \text{Mod}(\mathcal{K}_i^j)$ means that the pieces of information that agent i has about agent j 's preferences do not conflict with agent j 's real preferences.

Note that GPL, in contrast to e.g. CP-nets [Boutilier *et al.*, 2004], has the ability to model ignorance. Indeed, CP-nets can encode that an agent has no preference between two alternatives, but it cannot encode that we do not know whether the agent has a preference. On the other hand, GPL does assume that the preferences can be modeled by means of a total order. Importantly, GPL allows us to compactly describe information about another agent's preferences. For instance, through use of the operators \mathbf{N} and Δ , we can naturally encode necessary and sufficient conditions, respectively, for an agent to reach a certain utility. Moreover, we can easily encode *comparative preferences*, stating that the utility of an agent will always be higher for outcomes satisfying a formula α than for outcomes satisfying a formula β . To this end, we introduce some abbreviations for $\alpha, \beta \in L_{\Phi}$:

$$\begin{aligned} \beta \succeq \alpha &\equiv \bigvee_{m=1}^{p-1} (\neg \Pi_{\frac{m+1}{p}}(\alpha) \wedge \Delta_{\frac{m}{p}}(\beta)) \vee \neg \Pi_{\frac{1}{p}}(\alpha) \vee \Delta_1(\beta) \\ \beta \succ \alpha &\equiv \bigvee_{m=1}^p (\neg \Pi_{\frac{m}{p}}(\alpha) \wedge \Delta_{\frac{m}{p}}(\beta)) \end{aligned} \quad (2)$$

Intuitively, whenever $\mathcal{K}_i^j \models \alpha' \succeq \alpha$, agent i knows that the utility of agent j in any outcome that satisfies α' is at least the utility of j in any outcome that satisfies α . Similarly, whenever $\mathcal{K}_i^j \models \alpha' \succ \alpha$, agent i knows that agent j strictly prefers any outcome in which α' is true to any outcome in which α is true. Another useful abbreviation is:

$$\alpha' \succ_c \alpha \equiv \bigvee_{m=1}^p (\neg \Pi_{\frac{m}{p}}(\alpha) \wedge \Pi_{\frac{m}{p}}(\alpha'))$$

Intuitively, whenever $\mathcal{K}_i^j \models \alpha' \succ_c \alpha$, agent i knows that the outcome with the highest utility satisfying α' is strictly preferred to the outcome with the highest utility satisfying α . This allows us to model conditional preferences. For instance, $\alpha \wedge \beta_1 \succ_c \alpha \wedge \beta_2$ means that, in an ideal world, if α is true, the agent strictly prefers β_1 over β_2 . An exception x to the latter can be modeled by $\alpha \wedge x \wedge \beta_2 \succ_c \alpha \wedge x \wedge \beta_1$. This is similar in spirit to the possibilistic semantics of conditionals [Benferhat *et al.*, 1997]. Conditional preferences in possibilistic logic have been studied among others in [Dubois *et al.*, 2006].

We illustrate the expressiveness of GPL for modeling preferences in the following example.

Example 2

Recall the context of Example 1 and suppose Alice knows that Bob's first priority goal can only be fulfilled without bringing the dog. This is encoded as $\mathbf{N}_{\frac{1}{3}}(\neg d_A) \in \mathcal{K}_1^2$. If Alice knows that Bob prefers going to the beach exclusively with her over going to the beach with her and the dog, this

is encoded as $(b_A \wedge b_B \wedge \neg d_A) \succ (b_A \wedge b_B \wedge d_A) \in \mathcal{K}_1^2$. If Bob knows that Alice is unhappy without her dog, this is encoded as $\mathbf{N}_1(d_A) \in \mathcal{K}_2^1$. For Bob to encode that Alice is at least partially happy when she is with the dog, regardless of whatever else happens, he adds $\Delta_{\frac{1}{3}}(d_A)$ to \mathcal{K}_2^1 .

We now define a set of possibilistic discrimin optimal outcomes in \mathcal{V} . Intuitively, an outcome ν is optimal if for any outcome ν' which dominates ν according to the discrimin ordering, the agents who are better off in ν' than in ν are not aware that ν' is a valid counteroffer in the sense of (1).

Definition 7 (Possibilistic Discrimin)

We define the set of possibilistic discrimin optimal outcomes:

$$\begin{aligned} \text{Opt}_{discr}^p &= \{\nu \in \llbracket \delta \rrbracket \mid \forall \nu' \in \mathcal{V} : \nu' \succ_d \nu \Rightarrow (\forall i \in N : \\ &u_i(\nu') > u_i(\nu) \Rightarrow (\exists j \in N, \exists u_i^j \in \text{Mod}(\mathcal{K}_i^j) : \\ &u_i^j(\nu') < u_i^j(\nu) \wedge u_i^j(\nu') \leq u_i(\nu))\} \end{aligned}$$

It is easy to see that $\text{Opt}_{discr} \subseteq \text{Opt}_{discr}^p$. In particular, when each agent has full knowledge, i.e. $\text{Mod}(\mathcal{K}_i^j) = \{u_j\}$ for every $i, j \in N$, Opt_{discr} and Opt_{discr}^p coincide.

5 Negotiating under Incomplete Information

We now analyze negotiation in BGs under incomplete information. The protocol remains as specified in Algorithm 1: agents take turns in responding to an offer, by accepting it or making a counteroffer. However, the set of legal counteroffers $co(i, \nu)$ might be unknown to agent i . Indeed, determining the allowed counteroffers requires – possibly unknown – information about the other agents' utility. Therefore, we replace $co(i, \nu)$ by $co^p(i, \nu)$, which intuitively contains every outcome $\nu' \in \llbracket \delta \rrbracket$ for which agent i has enough information to derive that ν' is indeed a legal counteroffer to ν :

$$\begin{aligned} co^p(i, \nu) &= \{\nu' \in \llbracket \delta \rrbracket \mid (\mathcal{K}_i^i \models \nu' \succ \nu) \wedge \forall j \in N : \\ &\mathcal{K}_i^j \models (\Delta_{u_i(\nu) + \frac{1}{p}}(\varphi_{\nu'}) \vee \nu' \succeq \nu)\} \end{aligned}$$

with \succ and \succeq as defined in (2). As before, an outcome ν is agreed upon iff $co^p(i, \nu) = \emptyset$ for every $i \in N$.

We can prove that every possibilistic discrimin optimal outcome is generally accepted. To this end, we first prove the following link between the sets of valid counteroffers under complete and incomplete information.

Proposition 4

For every $i \in N$ and $\nu \in \mathcal{V}$: $co^p(i, \nu) \subseteq co(i, \nu)$.

Proof. Let ν' be an arbitrary element of $co^p(i, \nu)$. By definition, it holds that $u_i(\nu') > u_i(\nu) \wedge \forall j \in N : \mathcal{K}_i^j \models (\Delta_{u_i(\nu) + \frac{1}{p}}(\varphi_{\nu'}) \vee \nu' \succeq \nu)$. Let $j \in N$. It remains to prove that $u_j(\nu') < u_j(\nu) \Rightarrow u_j(\nu') > u_i(\nu)$ or equivalently $u_j(\nu') \geq u_j(\nu) \vee u_j(\nu') > u_i(\nu)$. By definition 6, it holds that u_j is a model of \mathcal{K}_i^j . Consequently, it holds that u_j satisfies either $\Delta_{u_i(\nu) + \frac{1}{p}}(\varphi_{\nu'})$ or $\nu' \succeq \nu$. In the first case, it holds that $u_j(\nu') \geq u_i(\nu) + \frac{1}{p}$ or thus $u_j(\nu') > u_i(\nu)$. In the second case, it holds that $u_j(\nu') \geq u_j(\nu)$. In any case it holds that $u_j(\nu') \geq u_j(\nu) \vee u_j(\nu') > u_i(\nu)$. \square

Proposition 5

For $\nu \in \text{Opt}_{discr}^p$ and $i \in N$ it holds that $co^p(i, \nu) = \emptyset$.

Proof. Let $\nu \in \text{Opt}_{discr}^p$ and $i \in N$. Suppose there exists some $\nu' \in co^p(i, \nu)$. If $\neg(\nu' >_d \nu)$, then Propositions 1 and 4 imply that $\nu' \notin co^p(i, \nu)$, a contradiction. Now assume that $\nu' >_d \nu$. Since $\nu' \in co^p(i, \nu)$ and $Mod(\mathcal{K}_i^i) = \{u_i\}$ it follows that $u_i(\nu') > u_i(\nu)$. Because $\nu \in \text{Opt}_{discr}^p$, we know that there exists some $j \in N$ and $u_i^j \in Mod(\mathcal{K}_i^j)$ such that $u_i^j(\nu') < u_i^j(\nu)$ and $u_i^j(\nu') \leq u_i(\nu)$. Consequently, u_i^j does not satisfy $\Delta_{u_i(\nu)+\frac{1}{p}}(\varphi_{\nu'})$, nor $\nu' \succeq \nu$. This contradicts the fact that $\mathcal{K}_i^j \models (\Delta_{u_i(\nu)+\frac{1}{p}}(\varphi_{\nu'}) \vee \nu' \succeq \nu)$. \square

Conversely, we can also show that only possibilistic discrimin optimal outcomes will be generally accepted.

Proposition 6

For $\nu \in \mathcal{V} \setminus \text{Opt}_{discr}^p$ there is an $i \in N$ with $co^p(i, \nu) \neq \emptyset$.

Proof. For $\nu \in \mathcal{V} \setminus \text{Opt}_{discr}^p$ there exists a $\nu' \in \mathcal{V}$ such that $\nu' >_d \nu$ and there exists an agent $i \in N$ such that $u_i(\nu') > u_i(\nu)$, and for every $j \in N$ and $u_i^j \in Mod(\mathcal{K}_i^j)$ it holds that $u_i^j(\nu') \geq u_i^j(\nu)$ or $u_i^j(\nu') > u_i(\nu)$. Since $Mod(\mathcal{K}_i^i) = \{u_i\}$, it follows that $\mathcal{K}_i^i \models \nu' \succ \nu$. Now let j be an arbitrary agent in N . For every model u_i^j of \mathcal{K}_i^j such that $u_i^j(\nu') \geq u_i^j(\nu)$ it holds that u_i^j models $\nu' \succeq \nu$. For every model u_i^j of \mathcal{K}_i^j such that $u_i^j(\nu') > u_i(\nu)$ it holds that u_i^j models $\Delta_{u_i(\nu)+\frac{1}{p}}(\varphi_{\nu'})$. Consequently, we have $\mathcal{K}_i^j \models (\Delta_{u_i(\nu)+\frac{1}{p}}(\varphi_{\nu'}) \vee \nu' \succeq \nu)$, and thus $\nu' \in co^p(i, \nu)$. \square

Note that since the number of possible offers is finite and because each offer must strictly improve the previous offer in terms of the discrimin ordering, the negotiation protocol always ends. From Corollary 5 and Proposition 6 we know that the possible outcomes at the end of the negotiation protocol are exactly the possibilistic discrimin optimal outcomes.

From $\text{Opt}_{discr} \subseteq \text{Opt}_{discr}^p$ and Proposition 5, it follows that any discrimin optimal offer is overall accepted under incomplete information. However, Example 3 shows that the opposite does not hold, i.e. a non-discrimin optimal outcome might be accepted under incomplete information.

Example 3

Suppose, in the context of Example 1, that Alice has absolutely no information concerning Bob's goals. If Bob may make the first offer and suggests to go to the beach together without the dog, Alice's utility is 0. Although this outcome is discrimin dominated by going to the beach with the dog, Alice is unable to make this counteroffer, because she does not know whether Bob's utility is at least 0.33 in that case or whether Bob's utility is at least the same as in his first offer.

Note that, in contrast to a fully informed agent, an agent with limited knowledge might not be able to open with a discrimin-optimal solution. It is clear that having no information leaves an agent in a very weak position. Indeed, if agent i knows nothing about the preferences of another agent, it holds that $co^p(i, \nu) = \emptyset$ for every $\nu \in \mathcal{V}$, thus agent i is obliged to

accept every offer. In contrast, an agent who has full knowledge knows all valid counteroffers and may be able to achieve a better outcome than in any discrimin optimal outcome, cfr. Bob in Example 3. Note that an agent with full knowledge can either use a safe or a risky selection function. Suppose for instance that there are only three possible utility vectors: (0.6, 0.4), (0.4, 0.6) and (1, 0.2). If agent 1 proposes (0.6, 0.4), it is certain that agent 2 accepts. Alternatively, if agent 1 proposes (1, 0.2) and agent 2 does not know that there exists a valid counteroffer, agent 1 can get away with an unfair agreement, yielding a higher utility than in any fair outcome. However, if agent 2 knows that (0.4, 0.6) is a valid counteroffer, the negotiations end in (0.4, 0.6), leaving agent 1 worse off than if it had proposed (0.6, 0.4) right away. This discussion shows that an interesting extension of the framework would be to allow agents to reason about the knowledge of others. Such knowledge can be encoded using multi-agent extensions of modal logics for epistemic reasoning, although we are then forced to express knowledge about preferences at the propositional level (e.g. by introducing variables g_i^m to denote the m^{th} most preferred goal of agent i as in [De Clercq et al., 2014]). This extension would allow agents to act based on their knowledge of how other agents would react to various counteroffers, as is common in the field of epistemic game theory [Perea, 2012].

6 Conclusion and Future Work

We have developed a framework for BGs with incomplete information, using GPL to compactly represent agents' knowledge about the preferences of others. We also proposed a multilateral negotiation protocol, which uses an intuitive negotiation rule based on the ethic of reciprocity principle and is guaranteed to converge within a finite number of steps. Moreover, we characterized the set of possible outcomes of the negotiation process, confirming the intuition that incomplete knowledge may lead to negotiation inefficiency, i.e. the agreement outcome may not be fair or efficient. In our protocol, the order of the agents plays an important role, which is natural in hierarchical contexts (e.g. leader-follower type setting, where followers can only question proposals by leaders if they can prove their unfairness). Alternatively, the power of agents [Ben-Naim and Lorini, 2014] can be used to deduce a sensible ordering in which agents are allowed to make offers: the most powerful agent can make the initial offer. Note, however, that the use of GPL for encoding knowledge about the preferences of others is independent of the negotiation protocol. Consequently, future research w.r.t. alternative negotiation protocols e.g. for settings in which agents have equal status can also rely on our GPL framework.

Even though the negotiation model we have discussed in this paper is rather simple, it offers a rich basis from which we can study a wide variety of settings. Interesting extensions could include the use of agents who expand their knowledge base during the protocol, by drawing conclusions from the offering behavior of other agents [Cramton, 1984]. Another option is to use different negotiation rules, e.g. an agent could be allowed to make a counteroffer ν if it does not know that ν is an illegal counteroffer. Additionally, we can allow

‘third party’ agents to protest against offers, in case they know that the offer is illegal. However, protesting against an unfair proposal requires the revelation of knowledge, which might weaken the bargaining power of the agent. Hence, it is not straightforward that protesting is always in the protester’s advantage, even if it initially leads to a higher utility. Other options for alternative protocols include the addition of time constraints [Kraus *et al.*, 1995] or the use of arguments to support an offer [Amgoud and Prade, 2003].

References

- [Ågotnes *et al.*, 2013a] T. Ågotnes, P. Harrenstein, W. van der Hoek, and M. Wooldridge. Boolean games with epistemic goals. In *Logic, Rationality, and Interaction*, volume 8196 of *LNCS*, pages 1–14. Springer, 2013.
- [Ågotnes *et al.*, 2013b] T. Ågotnes, P. Harrenstein, W. van der Hoek, and M. Wooldridge. Verifiable equilibria in Boolean games. In *Proc. IJCAI*, 2013.
- [Amgoud and Prade, 2003] L. Amgoud and H. Prade. A possibilistic logic modeling of autonomous agents negotiation. In *Progress in Artificial Intelligence*, pages 360–365. Springer, 2003.
- [Ben-Naim and Lorini, 2014] J. Ben-Naim and E. Lorini. Evaluating power of agents from dependence relations in Boolean games. In *Proc. AAMAS*, pages 853–860, 2014.
- [Benferhat *et al.*, 1997] S. Benferhat, D. Dubois, and H. Prade. Nonmonotonic reasoning, conditional objects and possibility theory. *Artif. Intell.*, 92(1):259–276, 1997.
- [Benferhat *et al.*, 2001] S. Benferhat, D. Dubois, and H. Prade. Towards a possibilistic logic handling of preferences. *Appl. Intell.*, 14(3):303–317, 2001.
- [Binmore *et al.*, 1986] K. Binmore, A. Rubinstein, and A. Wolinsky. The Nash bargaining solution in economic modelling. *The RAND Journal of Economics*, pages 176–188, 1986.
- [Bonzon *et al.*, 2006] E. Bonzon, M.-C. Lagasque-Schiex, and J. Lang. Compact preference representation for Boolean games. In *Proc. PRICAI*, pages 41–50. Springer-Verlag, 2006.
- [Bonzon *et al.*, 2012] E. Bonzon, M.-C. Lagasque-Schiex, and J. Lang. Effectivity functions and efficient coalitions in Boolean games. *Synthese*, 187:73–103, 2012.
- [Bossert and Peters, 2001] W. Bossert and H. Peters. Minimax regret and efficient bargaining under uncertainty. *Games and Economic Behavior*, 34(1):1–10, 2001.
- [Boutilier *et al.*, 2004] C. Boutilier, R.I. Brafman, C. Domshlak, H.H. Hoos, and D. Poole. CP-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *J. Artif. Intell. Res.*, 21:135–191, 2004.
- [Bouveret and Lang, 2008] S. Bouveret and J. Lang. Efficiency and envy-freeness in fair division of indivisible goods: Logical representation and complexity. *J. Artif. Intell. Res.*, pages 525–564, 2008.
- [Cramton, 1984] P.C. Cramton. Bargaining with incomplete information: An infinite-horizon model with two-sided uncertainty. *The Review of Economic Studies*, 51(4):579–593, 1984.
- [De Clercq *et al.*, 2014] S. De Clercq, S. Schockaert, M. De Cock, and A. Nowé. Possibilistic Boolean games: Strategic reasoning under incomplete information. In *Proc. JELIA*, pages 196–209, 2014.
- [Dubois *et al.*, 1994] D. Dubois, J. Lang, and H. Prade. Possibilistic logic. In *Handbook of Logic for Artificial Intelligence and Logic Programming*, volume 3, pages 439–513. Oxford University Press, 1994.
- [Dubois *et al.*, 1996] D. Dubois, H. Fargier, and H. Prade. Possibility theory in constraint satisfaction problems: Handling priority, preference and uncertainty. *Appl. Intell.*, 6(4):287–309, 1996.
- [Dubois *et al.*, 1997] D. Dubois, H. Fargier, and H. Prade. Beyond min aggregation in multicriteria decision: (ordered) weighted min, discrim, leximin. In *The ordered weighted averaging operators*, pages 181–192. Springer, 1997.
- [Dubois *et al.*, 2006] D. Dubois, S. Kaci, and H. Prade. Approximation of conditional preferences networks CP-nets in possibilistic logic. In *Proc. FUZZ-IEEE*, 2006.
- [Dubois *et al.*, 2014] D. Dubois, H. Prade, and S. Schockaert. Reasoning about uncertainty and explicit ignorance in generalized possibilistic logic. In *Proc. ECAI*, pages 261–266, 2014.
- [Dunne *et al.*, 2008] P.E. Dunne, W. van der Hoek, S. Kraus, and M. Wooldridge. Cooperative Boolean games. In *Proc. AAMAS*, volume 2, pages 1015–1022. IFAAMAS, 2008.
- [Fatima *et al.*, 2002] S.S. Fatima, M. Wooldridge, and N.R. Jennings. Optimal negotiation strategies for agents with incomplete information. In *Intelligent Agents VIII*, pages 377–392. Springer, 2002.
- [Grant *et al.*, 2011] J. Grant, S. Kraus, M. Wooldridge, and I. Zuckerman. Manipulating Boolean games through communication. In *Proc. IJCAI*, 2011.
- [Harrenstein *et al.*, 2001] P. Harrenstein, W. van der Hoek, J.-J. Meyer, and C. Witteveen. Boolean games. In *Proc. TARK*, pages 287–298. MKP Inc., 2001.
- [Harsanyi, 1967] J.C. Harsanyi. Games with incomplete information played by “Bayesian players”, I-III. *Management Science*, 14(3):159–182, 1967.
- [Hertzler, 1934] J.O. Hertzler. On golden rules. *Int. J. Ethics*, 44(4):418–436, 1934.
- [Kraus *et al.*, 1995] S. Kraus, J. Wilkenfeld, and G. Zlotkin. Multi-agent negotiation under time constraints. *Artif. Intell.*, 75(2):297–345, 1995.
- [Luo *et al.*, 2003] X. Luo, N.R. Jennings, N. Shadbolt, H. Leung, and J.H. Lee. A fuzzy constraint based model for bilateral, multi-issue negotiations in semi-competitive environments. *Artif. Intell.*, 148(1):53–102, 2003.
- [Marchioni and Wooldridge, 2014] E. Marchioni and M. Wooldridge. Łukasiewicz games. In *Proc. AAMAS*, pages 837–844, 2014.
- [Osborne and Rubinstein, 1994] M.J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT press, 1994.
- [Perea, 2012] A. Perea. *Epistemic game theory: reasoning and choice*. Cambridge University Press, 2012.
- [Tungodden, 2000] B. Tungodden. Egalitarianism: Is leximin the only option? *Econ. Philos.*, 16:229–245, 2000.