ECAI 2016
G.A. Kaminka et al. (Eds.)

doi:10.3233/978-1-61499-672-9-329

Formalizing Commitment-Based Deals in Boolean Games[1]

Sofie De Clercq[2] and **Steven Schockaert**[3] and **Ann Nowé**[4] and **Martine De Cock**[2,5]

Abstract. Boolean games (BGs) are a strategic framework in which agents' goals are described using propositional logic. Despite the popularity of BGs, the problem of how agents can coordinate with others to (at least partially) achieve their goals has hardly received any attention. However, negotiation protocols that have been developed outside the setting of BGs can be adopted for this purpose, provided that we can formalize (i) how agents can make commitments and (ii) how deals between coalitions of agents can be identified given a set of active commitments. In this paper, we focus on these two aims. First, we show how agents can formulate commitments that are in accordance with their goals, and what it means for the commitments of an agent to be consistent. Second, we formalize deals in terms of coalitions who can achieve their goals without help from others. We show that verifying the consistency of a set of commitments of one agent is Π_2^{P}-complete while checking the existence of a deal in a set of mutual commitments is Σ_2^{P}-complete. Finally, we illustrate how the introduced concepts of commitments and deals can be used to achieve game-theoretical properties of the deals and to configure negotiation protocols.

1 Introduction

Boolean games (BGs) are a game-theoretic framework which uses propositional logic to represent the goals of agents in a compact way [24]. A key feature of BGs is that each agent controls the truth value of a subset of the atoms from which these goals are built; these atoms are referred to as the action variables of the agent. In standard BGs, goals are of a binary nature [24]. In the context of negotiation, however, it is usually more natural to consider prioritized goal bases [6, 14], as these allow agents to partially concede. The basic intuition underlying BGs is illustrated in the next example.

Example 1
Suppose there are four agents, denoted by 1, 2, 3 and 4, representing four nations. Each agent i controls an action variable d_i. If agent i sets its action variable d_i to true, this means that i will disarm its nation. Nation 1 considers nation 2 a threat, nation 3 an ally and nation 4 irrelevant for its military strategy. It believes to be safe if either nation 2 disarms or nations 1 and 3 both keep their arms, i.e. $d_2 \vee (\neg d_1 \wedge \neg d_3)$. Nation 2 considers nation 1 as the only real threat, but prefers to disarm itself due to the associated costs of maintaining its arms. Therefore, its highest priority goal is $d_1 \wedge d_2$. Nation 3 strongly believes in the possibility of an alien invasion and prefers all nations to be armed, i.e. its highest priority goal is $\neg d_1 \wedge \neg d_2 \wedge \neg d_3 \wedge \neg d_4$. The pacifistic nation 4's first priority is the disarmament of all nations, i.e. $d_1 \wedge d_2 \wedge d_3 \wedge d_4$.

Although Boolean games have been widely studied in recent years, leading among others to the characterization of numerous solution concepts, the literature on BGs provides surprisingly few tools for agents to actually coordinate towards mutually beneficial agreements (see Section 5). The broader literature on multi-agent systems, however, has provided numerous negotiation protocols [34, 26, 3, 18, 33]. From a high-level point of view, many of these protocols are based on agents formulating commitments, intuitively encoding what they are prepared to offer in return for their goals being (partially) fulfilled. After a number of rounds, in which agents may progressively weaken their stance, the agents may end up with a set of mutual commitments which are such that a deal can be made. There are many technical details that need to be specified as part of a negotiation protocol (related e.g. to how agents communicate), but most of these are not dependent on how the agent's goals are encoded. In particular, to adapt existing negotiation protocols to the BG setting, it suffices to specify how agents can formulate commitments (i.e. proposal submission) and how deals based on these commitments can be made (i.e. agreement formation). The incorporation of the introduced notions into existing protocols is illustrated in Section 4.3.

Central to the discussion in this paper is the notion of a commitment. In the literature, a commitment is commonly stated as $(i; j; \mathit{ante}; \mathit{con})$, with the interpretation that agent i commits to agent j to bring about *con* when *ante* is made true by the other agents [37]. For instance, in the context of Example 1, a sensible commitment for agent 2 would be $(2; 1; d_1; d_2)$: if agent 1 disarms, agent 2 is prepared to do this as well. Commitments provide an intuitive way to formulate a propositional proposal and at the same time capture the fact that particular action variables are controlled by particular agents. This makes commitments a natural fit for the framework of BGs. Moreover, by identifying creditors, a commitment-based protocol allows the formation of deals between a subset of agents, i.e. the formation of coalitions. For instance, suppose in Example 1 that agent 1 and 2 respectively formulate the commitments $(1; 2; d_2; \top)$ and $(2; 1; d_1; d_2)$, where the former communicates agent 1's willingness to play any strategy if agent 2 disarms. Note such a commitment merely informs agent 2 of possibilities, but yields no guarantees, since $\top$ is brought about by default. The commitments of agents 1 and 2 lead to a possible deal: they can form a coalition $\{1, 2\}$ and play $\{d_1, d_2\}$, i.e. both nations disarm. To confirm the deal, however, agent 1 has to make a stronger commitment, i.e. it must specifically commit to bring about d_1 if d_2 is brought about. If this deal between agents 1 and 2 is closed, the BG can be reduced, allowing the remaining agents to update their goals: agent 3's

[1] This research was funded by a Research Foundation-Flanders project.
[2] Dept. of Applied Mathematics, CS and Statistics, Ghent University, Ghent, Belgium, {SofieR.DeClercq, Martine.DeCock}@ugent.be
[3] School of Computer Science and Informatics, Cardiff University, Cardiff, UK, SchockaertS1@cardiff.ac.uk
[4] Computational Modeling Lab, Vrije Universiteit Brussel, Brussels, Belgium, Ann.Nowe@vub.ac.be
[5] Center for Data Science, University of Washington, Tacoma, US, MDeCock@uw.edu

highest priority goal is no longer achievable and it should now turn to its lower priority goals. Agent 4's highest priority goal is reduced to $d_3 \wedge d_4$, i.e. in order to disarm all nations it remains to disarm itself and convince agent 3 to disarm as well. Note that by identifying deals between coalitions of agents, a global consensus is not required to obtain local deals. Moreover, the reduction of the BG induced by the previous deals facilitates matters for the agents that were not yet able to close a deal. These advantages are especially important in large-scale games, which many real-life applications are.

An important requirement for commitments is that agents should be able to guarantee that they can fulfill them. For instance, suppose that agent 1 in Example 1 communicates the commitments $(1; 2; d_2; d_1)$ and $(1; 3; \neg d_3; \neg d_1)$. Clearly, in case agent 2 and 3 play $\{d_2, \neg d_3\}$, the agent cannot play a strategy such that all its commitments are fulfilled. In Section 3 we formalize a notion of consistency, which captures the intuition that agents should not make commitments that cannot be jointly fulfilled. As we will show, checking whether a given set of commitments is consistent is a Π_2^P-complete problem.

Given a set of commitments, the main inference task we consider is verifying whether any corresponding deals can be made. Some of the issues underlying this process are illustrated in the next example.

Example 2
Suppose the agents in Example 1 communicate the following commitments:

$$(1; 2; d_2; \top) \qquad (3; \{1, 2, 4\}; \neg d_1 \wedge \neg d_2 \wedge \neg d_4; \neg d_3)$$
$$(1; 3; \neg d_3; \neg d_1) \qquad (4; \{1, 2, 3\}; d_1 \wedge d_2 \wedge d_3; d_4)$$
$$(2; 1; d_1; d_2)$$

Intuitively, the commitments agents 1 and 2 make to each other allow a deal between them to both disarm: $\{d_1, d_2\}$. However, as stated earlier, agent 1 has not specifically committed to bring about d_1 when d_2 is brought about. Therefore, we cannot consider this a confirmed deal. Moreover, agent 1's commitment to agent 3 blocks the possible deal with 2 when agent 3 decides to play $\neg d_3$. As a result, agents 1 and 2 cannot form a coalition by themselves based on the current set of commitments, as they would be reliant on what agent 3 subsequently decides.

Note that to identify deals, the control assignment of the BG should be taken into account, i.e. it does not suffice to check the satisfiability of conjunctions of formulas corresponding to goals of coalitions. Suppose, for instance, that agents 1 and 2 want agent 3 to disarm, i.e. the conjunction of their highest priority goal is satisfied if d_3 holds. Then no deal can be reached unless agent 3 can be convinced to set d_3 to true. As we will show, as a result, the problem of checking whether a deal can be made given a set of commitments is Σ_2^P-complete. To the best of our knowledge, this paper is the first to study commitment-based deals in BGs.

The paper is structured as follows. First, we give some background on BGs in Section 2. In Section 3 we formalize commitments in BGs, defining important concepts, including the consistency of a set of commitments, which guarantees that agents can honour their commitments, irrespective of the strategies of the other agents. After investigating the computational complexity of verifying consistency in Section 3.1, we explain how an agent can formulate commitments that accord with a single goal or with a prioritized goal base in respectively Sections 3.2 and 3.3. Then we formalize how agents can identify deals based on a set of commitments and investigate the computational complexity in Section 4.1. Next, we illustrate how our concepts can be used to guarantee game-theoretical properties of the deals in Section 4.2 and how they can be implemented in existing protocols in Section 4.3. We discuss related work in Section 5 and present our conclusion and interesting questions for further research in Section 6. The proofs of all results are available in an online appendix[6].

2 Preliminaries

The propositional language L_Φ is built from a finite set of atoms Φ in the usual way. We write $\mathit{Lit}(\Phi) = \Phi \cup \{\neg p \mid p \in \Phi\}$. An interpretation of Φ is defined as a subset ν of $\mathit{Lit}(\Phi)$ such that for every atom $p \in \Phi$ either $p \in \nu$ or $\neg p \in \nu$. We denote the set of all interpretations of Φ as $\mathit{Int}(\Phi)$. An interpretation can be extended to L_Φ in the usual way. We write $\nu \models \varphi$ to denote that formula φ is true in interpretation ν. Whenever we write an interpretation ν where a formula is expected, this should be interpreted as the conjunction of ν's literals. For two formulas φ and $\psi \in L_\Phi$, it holds that φ entails ψ, denoted $\varphi \models \psi$, iff for every interpretation ν it holds that $\nu \models \psi$ whenever $\nu \models \varphi$. We say that a variable p is irrelevant in a formula if there exists an equivalent formula in which p does not occur [30]. The relevant variables of a formula γ are denoted as $\mathit{DepVar}(\gamma)$. We say that φ and ψ are *equivalent modulo* δ, denoted $\varphi \equiv \psi (\mathrm{mod}\, \delta)$, iff $\varphi \wedge \delta$ is equivalent with $\psi \wedge \delta$ [30]. We use the variant of Boolean games from [14].

Definition 1 (Boolean Game with Priorities)
A Boolean game (BG) with priorities is a tuple $G = (N, (\Phi_i)_{i \in N}, (\delta_i)_{i \in N}, (\Gamma_i)_{i \in N})$, where N is a finite set of agents, Φ_i is a finite set of atoms such that $\Phi_i \cap \Phi_j = \emptyset$ for $j \neq i$, δ_i is a consistent formula in L_{Φ_i}, and $\Gamma_i = \{\gamma_i^1; \ldots; \gamma_i^p\}$ is i's prioritized goal base. We write $\Phi = \bigcup_{i \in N} \Phi_i$ and $\delta = \bigwedge_{i \in N} \delta_i$. The formula $\gamma_i^m \in L_\Phi$ is agent i's goal of priority m. We assume w.l.o.g. that every agent has p priority levels.

Definition 1 is a particular case of generalized BGs [6] in which the preference relations are total, but with the addition of constraints δ_i [8]. The set Φ contains all action variables (or atoms). Agent i controls the truth value of the atoms in Φ_i, with the restriction that δ_i must be satisfied. We write $\mathcal{S}_i = \{\nu_i \in \mathit{Int}(\Phi_i) \mid \nu_i \models \delta_i\}$ for the set of permissible strategies of agent i. A non-empty subset of N is called a coalition. We straightforwardly extend definitions w.r.t. agents to coalitions, e.g. Φ_J is the set of action variables controlled by a coalition J, $\mathcal{S}_J$ denotes the permissible joint strategies. We denote singleton coalitions $\{i\}$ as i when there can be no confusion. By convention, goals are ordered from high (level 1) to low priority (level p).

Example 1 (Continued)
We have $N = \{1, 2, 3, 4\}$, $\Phi_i = \{d_i\}$ and $\delta_i = \top$ for each $i \in N$. Consequently, $\mathcal{S}_i = \{d_i, \neg d_i\}$, i.e. each agent can either disarm or not. The prioritized goal bases in the BG are:

$$\Gamma_1 = \{d_2 \vee (\neg d_1 \wedge \neg d_3); \neg d_1\}, \Gamma_2 = \{d_1 \wedge d_2; \neg d_2\}$$
$$\Gamma_3 = \{\neg d_1 \wedge \neg d_2 \wedge \neg d_3 \wedge \neg d_4; (\neg d_1 \vee \neg d_2 \vee \neg d_4) \wedge \neg d_3\}$$
$$\Gamma_4 = \{d_1 \wedge d_2 \wedge d_3 \wedge d_4; d_4\}$$

This means, e.g. for agent 2, that its first priority is for agents 1 and 2 to disarm and its second priority is to arm.

We now define the concept of relevant agents for a formula, which is related to the concept of relevant agents for another agent described in [7].

[6] http://www.cwi.ugent.be/sofie/ECAI16appendix.pdf

Definition 2 (Relevant Agents for a Formula)
The set of relevant agents $RelAg(\varphi)$ *for* $\varphi \in L_\Phi$ *is defined as*

$$\bigcup_{i\in N}\Big\{\Big(\bigcap_{\psi\equiv\varphi(mod\,\delta)} DepVar(\psi)\Big)\cap\Phi_i \neq \emptyset\Big\}$$

In other words, the relevant agents for a formula are the agents controlling its relevant variables, taking into account the constraint δ. For instance, in Example 1, the relevant agents for $(d_1 \wedge d_2)\wedge(\neg d_3 \vee d_3)$ are $\{1, 2\}$.

Definition 3 (Outcome)
An interpretation of Φ *satisfying* δ *is called an outcome of* G*. We denote the set of all outcomes as* $\mathcal{S}_N$.

An outcome ν corresponds to a tuple $(\nu_1, \dots, \nu_n)$ with $\nu_i \in \mathcal{S}_i$ for every $i \in N$. We write $\nu_J = \nu \cap Lit(\Phi_J)$ for the restriction of ν to the strategies of agents in coalition J. The restriction of ν to the agents outside J is denoted as ν_{-J}. For disjoint coalitions J and J', $\nu \in \mathcal{S}_J$ and $\nu' \in \mathcal{S}_{J'}$, we write $(\nu, \nu') \in \mathcal{S}_{J\cup J'}$ to denote their joint strategy. We also use the notation (ν, ν') if J and J' are not disjoint, but $\nu_i = \nu'_i$ for every $i \in J \cap J'$. For the ease of presentation, we define a numerical utility in $[0, 1]$, which is determined by the highest priority for which the corresponding goal is satisfied by ν.

Definition 4 (Utility Function)
Let G *be a BG. For each* $i \in N$ *and* $\nu \in \mathcal{S}_N$*, the utility of* i *in* ν *is defined as:*

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

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

Note that if ν does not satisfy any goal, the utility is 0, while it is 1 iff the first priority goal is satisfied. In Example 1, for instance, the utility of agent 1 is 1 for every outcome ν in which the coalition $\{1, 3\}$ plays the joint strategy $\{\neg d_1, \neg d_3\}$, i.e. if both nations 1 and 3 decide against disarmament. Note that there exist alternative ways to extract utilities from prioritized goal bases [6], for which similar results as the ones presented in this paper can be obtained.

A well-known solution concept in BGs is the pure Nash equilibrium. This notion is based on best responses: we say that $\nu \in \mathcal{S}_i$ is a *best response* to $\nu^* \in \mathcal{S}_{-i}$, written $\mathcal{BR}(G, \nu, \nu^*)$, iff for every $\nu' \in \mathcal{S}_i$ it holds that $u_i(G, (\nu, \nu^*)) \geq u_i(G, (\nu', \nu^*))$. Intuitively, this means that given the strategies of the other agents, an agent is not better off by deviating from its current strategy. In Example 1, for instance, agent 2's first priority is the disarmament of agents 1 and 2. Therefore, if agent 1 disarms, agent 2's best response is to disarm as well.

Definition 5 (Pure Nash Equilibrium)
An outcome ν *is a pure Nash equilibrium (PNE) of the BG* G *iff* $\mathcal{BR}(G, \nu_i, \nu_{-i})$ *holds for every* $i \in N$.

In Example 1, for instance, the outcomes $\{d_1, d_2, \neg d_3, d_4\}$ and $\{\neg d_1, \neg d_2, \neg d_3, d_4\}$ are PNEs.

3 Formalizing Commitments in BGs

In this section, we adapt the notion of commitment from [37] to the context of BGs. Among others, we analyze how agents can formalize consistent commitments based on the goal they want to achieve. Throughout this section, we will assume that $G = (N, (\Phi_i)_{i\in N}, (\delta_i)_{i\in N}, (\Gamma_i)_{i\in N})$ is a BG.

Definition 6 (Commitment in BG)
The tuple $c = (deb(c); cred(c); ante(c); con(c))$ *is called a commitment in* G *if* $deb(c) \in N$*,* $cred(c) \subseteq N \setminus \{deb(c)\}$*,* $ante(c)$ *is a formula such that* $RelAg(ante(c)) = cred(c)$*,* $ante(c) \wedge \delta \not\models \bot$*,* $con(c)$ *is a formula containing only variables from* $\Phi_{deb(c)}$*, and* $con(c) \wedge \delta \not\models \bot$.

Intuitively, a commitment c describes a state of affairs in which the debtor $deb(c)$ commits to the creditors[7] $cred(c)$ to bring about the consequent $con(c)$ if the antecedent $ante(c)$ is satisfied. To exclude meaningless commitments, we assume that $con(c)$ contains only variables from $\Phi_{deb(c)}$, i.e. we do not allow an agent to commit to anything outside its own control. Moreover, we assume that $con(c)$ is consistent with $\delta_{deb(c)}$, meaning that there exists at least one strategy $\nu \in \mathcal{S}_{deb(c)}$ such that $\nu \models con(c)$. This restriction, which excludes impossible promises, corresponds to the *consistency postulate* for active commitments in [37]. We also exclude the debtor as one of its own creditors, as the commitment is meaningless to the other creditors if the debtor itself can control whether or not the antecedent is fulfilled. The condition $RelAg(ante(c)) = cred(c)$ makes the formulation of other definitions more convenient, but is not an explicit requirement for commitments [37]. We demand that $ante(c)$ is satisfiable w.r.t. δ, in line with the *nonvacuity postulate* for active commitments [37].

If $ante(c) \equiv \top$ (and thus $cred(c) = \emptyset$), the commitment is called unconditional. The case where $con(c) \equiv \top$ contradicts the first postulate in [37], which says that a commitment is *discharged* (no longer active) when the consequent holds. This makes sense, since such a commitment makes no guarantees to the creditor. However, as illustrated in Section 1, we will use such commitments to allow an agent to express that it is prepared to bring about any of its strategies when the antecedent is satisfied. This will be useful in our formalization of a possible deal, as will become clear in Definition 12. To obtain a confirmed deal though, the agent will have to make a stronger commitment, as will become clear in Definition 14.

In the following example and throughout this paper, subscripts are used to indicate the controlling agent of each action variable.

Example 3
In the context of Example 1, consider the pair of commitments $(1; 2; d_2; d_1)$ and $(1; 3; \neg d_3; \neg d_1)$. In the first commitment, agent 1 commits to disarm, i.e. bring about d_1, when agent 2 disarms, i.e. brings about d_2. In the second commitment, agent 1 commits to arm, i.e. bring about $\neg d_1$, when agent 3 arms, i.e. brings about $\neg d_3$.

Note that the commitments in Example 3 cannot be fulfilled simultaneously. We will call such commitments inconsistent and formalize the concept of consistent commitments in the following section.

3.1 Consistency of Commitments

Given that commitments are meant to be binding, it is important that an agent can jointly fulfil the set of all commitments it has made. To formalize this notion of consistency, we introduce the following abbreviations:

$$deb(\mathcal{C}) = \bigcup_{c\in\mathcal{C}} deb(c), \qquad cred(\mathcal{C}) = \bigcup_{c\in\mathcal{C}} cred(c),$$
$$ante(\mathcal{C}) = \bigwedge_{c\in\mathcal{C}} ante(c), \qquad con(\mathcal{C}) = \bigwedge_{c\in\mathcal{C}} con(c).$$

[7] In contrast to [37], we allow multiple creditors.

Definition 7 (Consistent Commitments)
Suppose $\mathcal{C}$ is a set of commitments of agent i, i.e. $deb(\mathcal{C}) = i$. Then $\mathcal{C}$ is consistent iff for every non-empty subset $\mathcal{C}'$ of $\mathcal{C}$:

$$(ante(\mathcal{C}') \wedge \delta) \text{ is consistent } \Rightarrow (con(\mathcal{C}') \wedge \delta) \text{ is consistent}$$

Intuitively, whenever there exists a strategy of the creditors of a subset $\mathcal{C}'$ of $\mathcal{C}$ that satisfies the antecedents in $\mathcal{C}'$, there should exist a strategy of the debtor of $\mathcal{C}'$ that satisfies all consequents in $\mathcal{C}'$. Note that it is not sufficient for consistency that every pair of commitments in the set $\mathcal{C}$ is consistent, as illustrated in the following example.

Example 4
Every pair in the following set of commitments is consistent, but the three commitments together are not: $(1; 2; d_2; d_1 \leftrightarrow \neg b_1)$, $(1; 3; d_3; b_1 \leftrightarrow \neg c_1)$, $(1; 4; d_4; c_1 \leftrightarrow \neg d_1)$.

Proposition 1
Deciding whether a given set of commitments with fixed debtor is consistent, is Π_2^P-complete.

3.2 Relating Goals and Commitments

In this section, we investigate how an agent with a goal γ can formulate a corresponding commitment.

Example 5
Reconsider the context of Example 1. Let $\gamma = d_1 \wedge d_2 \wedge d_3 \wedge d_4$ be a goal of agent 4. To achieve this goal, agent 4 can make the commitment $c = (4; \{1, 2, 3\}; d_1 \wedge d_2 \wedge d_3; d_1)$, which intuitively expresses that agent 4 commits to play a strategy that involves its disarmament if agents 1, 2 and 3 do the same.

We now formalize how a commitment can be created to match a propositional goal.

Definition 8 (Commitments Coinciding with Goal)
The set $\mathcal{C}$ of commitments with debtor i coincides with the goal γ of agent i iff for every $\nu \in \mathcal{S}_N$:

$$(\nu \models \gamma) \Leftrightarrow \big(\big(\exists c \in \mathcal{C} : (\nu_{-i} \models ante(c)) \wedge (\nu_i \models con(c))\big) \\ \wedge \big(\nu_i \models con(\{c \in \mathcal{C} \mid \nu_{-i} \models ante(c)\})\big)\big) \quad (1)$$

The '$\Rightarrow$' direction expresses that the set of commitments *covers* the goal γ, i.e. for every outcome ν that satisfies γ, i has an active commitment c corresponding to ν, which does not result in the violation of any other commitment of i. The '$\Leftarrow$' direction expresses that the set of commitments *is covered* by the goal γ, i.e. any outcome made possible by a commitment of i results in the satisfaction of γ.

Proposition 2
For any formula $\gamma \in L_\Phi$ and agent $i \in N$, there exists a consistent set $\mathcal{C}$ of commitments with debtor i that coincides with γ. Moreover, if γ can be rewritten as a conjunction of literals, then $\mathcal{C}$ can be chosen as a singleton.

It is easy to verify that some goals cannot coincide with a singleton, e.g. the goal $(d_1 \leftrightarrow d_2)$ requires two separate commitments. We illustrate the specification of commitments coinciding with goals in the following example.

Example 6
Suppose that $\gamma = (d_1 \wedge d_2) \vee (\neg d_1 \wedge \neg d_3)$ is a goal of agent 1. The set $\mathcal{C} = \{(1; \{2, 3\}; d_2 \wedge d_3; d_1), (1; \{2, 3\}; \neg d_2 \wedge \neg d_3; \neg d_1),$ $(1; \{2, 3\}; d_2 \wedge \neg d_3; \top)\}$ coincides with γ.

3.3 Commitments for Prioritized Goal Bases

We now relate prioritized goal bases and commitments.

Definition 9 (Commitments Guaranteeing Utility)
Let $\Gamma = \{\gamma^1; \ldots; \gamma^p\}$ be a prioritized goal base of i. The set $\mathcal{C}$ of commitments with debtor i guarantees utility k iff for each $\nu \in \mathcal{S}_N$:

$$u_i(G, \nu) \geq k \Leftarrow \big(\big(\nu_i \models con(\{c \in \mathcal{C} \mid \nu_{-i} \models ante(c)\})\big) \\ \wedge \big(\exists c \in \mathcal{C} : (\nu_{-i} \models ante(c)) \wedge (\nu_i \models con(c))\big)\big) \quad (2)$$

A straightforward way to construct a set of commitments that guarantees utility k is by constructing the set of commitments that coincides with formula $\bigvee_{m=1}^{p-kp+1} \gamma^m$, using the construction from Proposition 2. Note that for this particular choice, the '$\Rightarrow$' direction of (2) also holds. However, we do not require this direction in Definition 9, due to incompatibility with Definition 10, which formalizes when an agent's commitments are in line with its best responses.

Definition 10 (Commitment Respecting Best Responses)
The commitment c with debtor i respects i's best responses in G iff for each $\nu \in \mathcal{S}_N$:

$$\big((\nu_{-i} \models ante(c)) \wedge (\nu_i \models con(c))\big) \Rightarrow \mathcal{BR}(G, \nu_i, \nu_{-i}) \quad (3)$$

A set of commitments with debtor i respects i's best responses iff every commitment in the set does.

Example 7
Reconsider the context of Example 1. Suppose the pacifistic nation 4 has the goal base $\Gamma_4 = \{d_1 \wedge d_2 \wedge d_3 \wedge d_4; \neg d_4\}$, i.e. the first priority of nation 4 is global disarmament, but in case this turns out to be unachievable, it distrusts the other nations and prefers to maintain its arms. The commitment $(4; \emptyset; \top; \neg d_4)$ guarantees utility 0.5, but does not respect 4's best responses: if the other nations decide to disarm, 4 would have been better off by disarming as well. The commitment $(4; \{1, 2, 3\}; d_1 \wedge d_2 \wedge d_3; d_4)$ guarantees utility 1 and respects agent 4's best responses. The pair of commitments $(4; \{1, 2, 3\}; d_1 \wedge d_2 \wedge d_3; d_4)$ and $(4; \{1, 2, 3\}; \neg d_1 \vee \neg d_2 \vee \neg d_3; \neg d_4)$ guarantees utility 0.5 and respects agent 4's best responses.

Note that a set $\mathcal{C}$ of commitments with debtor i which respect i's best responses is not necessarily consistent. Consider for instance a BG with $\Gamma_1 = \{(a_1 \vee b_1) \wedge a_2\}$ and $\delta_1 = \neg(a_1 \wedge b_1)$. Agent 1's commitments $(1; 2; a_2; a_1)$ and $(1; 2; a_2; b_1)$ both respect its best responses, but are not consistent. However, we can show the following result.

Proposition 3
Let $\Gamma = \{\gamma^1; \ldots; \gamma^p\}$ be the goal base of some agent i. For any $k \in \{\frac{1}{p}, \ldots, \frac{p}{p}\}$ such that $\bigvee_{m=1}^{p-kp+1} \gamma^m \wedge \delta \not\models \perp$ there exists a non-empty consistent set $\mathcal{C}$ of commitments with debtor i that guarantees utility k and respects i's best responses.

4 Commitment-based Deals in BGs

From a high-level point of view, many negotiation protocols are based on agents formulating proposals or commitments, intuitively encoding what they are prepared to offer in return for their goals being (partially) fulfilled. After a number of rounds, in which agents may progressively weaken their stance, the agents may end up with a set of mutual commitments which are such that a deal can be made. Note that a deal does not require the involvement of all agents. For instance, suppose two neighbouring nations 1 and 2's first priority it to both disarm. In order to play the coalition strategy $\{d_1, d_2\}$, they do not need the approval of all other nations, as the action variables involved in the deal are controlled by the agents that closed the deal.

As deals can be closed between coalitions of agents, the BG can iteratively be reduced based on the chosen strategies of the agents who have already closed a deal. To formalize this, we use the notion of a formula $\varphi \in L_\Phi$ being conditioned by an interpretation ν of $\Phi' \subseteq \Phi$ [13], written as $cond(\varphi, \nu)$. The formula $cond(\varphi, \nu) \in L_{\Phi \setminus \Phi'}$ is obtained from φ in the following way: for every atom $p \in \Phi'$ such that $\neg p \in \nu$ we replace every occurrence of p in φ by $\bot$, and for every atom $p \in \Phi'$ such that $p \in \nu$ we replace every occurrence of p in φ by $\top$. Next, the tautologies $\neg\top \equiv \bot$, $\neg\bot \equiv \top$, $(\top \wedge \varphi) \equiv \varphi$, $(\top \vee \varphi) \equiv \top$, $(\bot \wedge \varphi) \equiv \bot$ and $(\bot \vee \varphi) \equiv \varphi$ are iteratively used to simplify the formula.

Definition 11 (Reduced BG)
Let $G = (N, (\Phi_i)_{i\in N}, (\delta_i)_{i\in N}, (\Gamma_i)_{i\in N})$, J a coalition of N and $\nu \in \mathcal{S}_J$. The reduced BG is defined as $red(G, J, \nu) = (N \setminus J, (\Phi_i)_{i\in N\setminus J}, (\delta_i)_{i\in N\setminus J}, (cond(\Gamma_i, \nu))_{i\in N\setminus J})$, with $cond(\Gamma_i, \nu) = \{cond(\gamma_i^1, \nu), \ldots, cond(\gamma_i^p, \nu)\}$ for $i \in N \setminus J$.

The next example illustrates this concept, which generalizes the notion of a *projection* from [7].

Example 8
Reconsider the context of Example 1, i.e. the BG G with four agents. Each agent i controls one variable, i.e. $\Phi_i = \{d_i\}$. Consider the coalition $\{1, 2\}$ with the joint strategy to disarm, i.e. $\{d_1, d_2\}$. Then the reduced game $red(G, \{1, 2\}, \{d_1, d_2\})$ consists of two agents (namely agents 3 and 4), with Φ_3 and Φ_4 as in G. The goal bases of agents 3 and 4 are reduced from

$$\Gamma_3 = \{\neg d_1 \wedge \neg d_2 \wedge \neg d_3 \wedge \neg d_4; (\neg d_1 \vee \neg d_2 \vee \neg d_4) \wedge \neg d_3\}$$
$$\Gamma_4 = \{d_1 \wedge d_2 \wedge d_3 \wedge d_4; d_4\}$$

to

$$\Gamma_3 = \{\bot; \neg d_4 \wedge \neg d_3\} \qquad \Gamma_4 = \{d_3 \wedge d_4; d_4\}$$

Note that in particular, agent 3 can no longer achieve utility 1.

The following results are straightforward to prove.

Proposition 4
Let G be a BG with J and J' two disjoint coalitions, $\nu \in \mathcal{S}_J$ and $\nu' \in \mathcal{S}_{J'}$. It holds that $red(red(G, J, \nu), J', \nu') = red(G, J \cup J', (\nu, \nu'))$.

Proposition 5
Let G be a BG, J a coalition and $\nu \in \mathcal{S}_J$. For $i \in N \setminus J$ and $\nu' \in \mathcal{S}_{N\setminus J}$ it holds that $u_i(G, (\nu, \nu')) = u_i(red(G, J, \nu), \nu')$.

To illustrate Proposition 5, note that in Example 8 agent 4's utility in $\{d_1, d_2, d_3, d_4\}$ in G is 1, which is the same as its utility in $\{d_3, d_4\}$ in the reduced game $red(G, \{1, 2\}, \{d_1, d_2\})$. The following corollary follows immediately from Proposition 5.

Corollary 6
Let G be a BG, J a coalition and $\nu \in \mathcal{S}_J$. For $i \in N \setminus J$, $\nu' \in \mathcal{S}_i$ and $\nu^* \in \mathcal{S}_{N\setminus(J\cup\{i\})}$ it holds that $\mathcal{BR}(G, \nu', (\nu, \nu^*)) \Leftrightarrow \mathcal{BR}(red(G, J, \nu), \nu', \nu^*)$.

Intuitively, Corollary 6 expresses that the reduced game preserves the best responses of the original game. For instance, in Example 8 the strategy $\{d_4\}$ is agent 4's unique best response to $\{d_3\}$ in $red(G, \{1, 2\}, \{d_1, d_2\})$, as well as its unique best response to $\{d_1, d_2, d_3\}$ in the original game G. As a consequence, PNEs are also preserved.

Corollary 7
Let G be a BG with $J \subset N$, ν a strategy of J and ν' a strategy of $N \setminus J$. It holds that (ν, ν') is a PNE of G iff ν' is a PNE of $red(G, J, \nu)$ and for every $i \in J$ it holds that $\mathcal{BR}(G, \nu_i, (\nu_{-i}, \nu'))$.

For instance, in Example 8 the outcome $\{d_1, d_2, d_3, d_4\}$ is a PNE of G. It also holds that $\{d_3, d_4\}$ is a PNE of the reduced game $red(G, \{1, 2\}, \{d_1, d_2\})$ and d_1 and d_2 are best responses to $\{d_2, d_3, d_4\}$ respectively $\{d_1, d_3, d_4\}$ in G. Similarly, $\{d_1, d_2\}$ is a PNE of $red(G, \{3, 4\}, \{d_3, d_4\})$ and d_3 and d_4 are best response to $\{d_1, d_2, d_4\}$ respectively $\{d_1, d_2, d_3\}$ in G.

Now that we have explained how deals between coalitions can be used to reduce the BG, it remains to formalize how agents can identify these deals based on a given set of commitments.

4.1 Identifying Deals

Given a set of commitments, intuitively, a possible deal corresponds to a coalition with a joint strategy such that all agents in the coalition actively support the coalition strategy through their commitments. Throughout this paper we assume that commitments are either common knowledge or known by one central entity.

Definition 12 (Possible Deal)
Let $\mathcal{C} = (\mathcal{C}_i)_{i\in N}$ be a tuple with $\mathcal{C}_i$ a set of commitments for every agent $i \in N$. The coalition J and strategy $\nu \in \mathcal{S}_J$ correspond to a possible deal based on $\mathcal{C}$ iff:

$$\forall i \in J, \exists c \in \mathcal{C}_i : ((\nu_{-i} \wedge \delta) \models ante(c)) \wedge (\nu_i \models con(c)) \tag{4}$$

Intuitively, Definition 12 expresses that a deal between a coalition of agents must be backed up by an active commitment of every participating agent, i.e. no agent agrees to a deal without benefiting from it. The computational complexity of the associated decision problem is situated at the second level of the polynomial hierarchy.

Proposition 8
Deciding whether there exists a possible deal based on a given set $\mathcal{C}$ of commitments is Σ_2^P-complete.

Condition (4) is obviously a necessary condition to reach an agreement. However, participating in a possible deal may require the agent to play a strategy which is incompatible with some of its other commitments. Moreover, a commitment might make a deal possible, yet not be strong enough to guarantee it. We illustrate these issues with an example.

Example 9
Reconsider the context of Example 1 and Example 2: a BG with $N = \{1, 2, 3, 4\}$, $\Phi_i = \{d_i\}$, $\delta = \top$, $\Gamma_1 = \{(\neg d_1 \wedge \neg d_3) \vee d_2; \ldots\}$, $\Gamma_2 = \{d_1 \wedge d_2; \ldots\}$, $\Gamma_3 = \{\neg d_1 \wedge \neg d_2 \wedge \neg d_3 \wedge \neg d_4; \ldots\}$ and $\Gamma_4 = \{d_1 \wedge d_2 \wedge d_3 \wedge d_4; \ldots\}$. Suppose the agents announce the following commitments:

$$\mathcal{C}_1 = \{c_1 = (1; 3; \neg d_3; \neg d_1), c_1' = (1; 2; d_2; \top)\}$$
$$\mathcal{C}_2 = \{c_2 = (2; 1; d_1; d_2)\}$$
$$\mathcal{C}_3 = \{c_3 = (3; \{1, 2, 4\}; \neg d_1 \wedge \neg d_2 \wedge \neg d_4; \neg d_3)\}$$
$$\mathcal{C}_4 = \{c_4 = (4; \{1, 2, 3\}; d_1 \wedge d_2 \wedge d_3; d_4)\}$$

These sets are consistent and respect each agent's best responses. The unique possible deal is $(\{1, 2\}, \{d_1, d_2\})$: c_1' and c_2 back up this deal. However, if agent 3 decides to bring about $\neg d_3$, agent 1 cannot play $\{d_1\}$ without violating its commitment c_1. Moreover, agent 1 has not specifically committed to bring about d_1 when d_2 is brought about.

To address these issues, we introduce the notion of *confirmed deals*. To this end, we first define the concept of a stable set of commitments, which captures the intuition of a coalition whose commitments entail the willingness of all members to participate in a corresponding deal, i.e. to jointly satisfy the antecedents of the commitments that are part of the deal.

Definition 13 (Stable Set of Commitments)
A set of commitments $\mathcal{C}$ with $J = deb(\mathcal{C})$ is called stable iff $con(\mathcal{C}) \wedge \delta \models ante(\mathcal{C})$ and $\mathcal{C}$ has at least one playable coalition strategy, i.e. $\exists \nu \in \mathcal{S}_J$ such that $\nu \models con(\mathcal{C})$.

In Example 9 there are no stable sets of commitments. If agent 1 replaces the commitment c_1 with $(1; \{2,3\}; \neg d_3 \wedge \neg d_2; \neg d_1)$ and c_1' with $c_1'' = (1; 2; d_2; d_1)$, then $\{c_1'', c_2\}$ is a stable set of commitments with $\{d_1, d_2\}$ its unique playable coalition strategy.
As the consequent of any commitment c contains only variables of $\Phi_{deb(c)}$ — see Definition 6 — it follows that $\nu \models con(\mathcal{C})$ is equivalent to $\nu_i \models con(\{c \in \mathcal{C} \,|\, deb(c) = i\})$ for every $i \in deb(\mathcal{C})$. Therefore, one can unambiguously speak of a playable strategy of an agent in $deb(\mathcal{C})$.

Definition 14 (Confirmed Deal)
Let $\mathcal{C} = (\mathcal{C}_i)_{i \in N}$ be such that each $\mathcal{C}_i$ only contains commitments with debtor i and let $\mathcal{C}' \subseteq \mathcal{C}$ be a stable set of commitments with $J = deb(\mathcal{C}')$. Then J and $\mathcal{C}'$ form a confirmed deal based on $\mathcal{C}$ iff $\mathcal{C}'$ has at least one safely playable coalition strategy, i.e. there is some $\nu' \in \mathcal{S}_J$ such that $\nu' \models con(\mathcal{C}')$, and for every $i \in J$ and $c \in \mathcal{C}_i$ we have

$$\forall \nu \in \mathcal{S}_{N \setminus J} : ((\nu, \nu'_{-i}) \models ante(c)) \Rightarrow (\nu'_i \models con(c)) \qquad (5)$$

For $J = N$, we drop the universal quantifier and ν in (5).

Intuitively, a confirmed deal guarantees for all coalition partners that the antecedent of their commitments in $\mathcal{C}'$ can be satisfied and that they can at the same time honour all their commitments in $\mathcal{C}$, regardless of what the agents outside the coalition do. This is achieved by playing a safely playable coalition strategy. The computational complexity of the associated decision problem is also situated at the second level of the polynomial hierarchy.

Proposition 9
Deciding whether there exists a confirmed deal based on a given set $\mathcal{C}$ of commitments is Σ_2^P-complete.

It turns out that if the coalition partners have expressed consistent commitments, having a stable set is sufficient to obtain a confirmed deal.

Proposition 10
Let $\mathcal{C} = (\mathcal{C}_i)_{i \in N}$ be such that each $\mathcal{C}_i$ is a consistent set of commitments with debtor i. If $\mathcal{C}' \subseteq \mathcal{C}$ is a stable set of commitments, then $(deb(\mathcal{C}'), \mathcal{C}')$ is a confirmed deal based on $\mathcal{C}$.

Moreover, a connection can be made between stable sets and possible deals.

Proposition 11
Let $\mathcal{C}$ be a stable set of commitments with $J = cred(\mathcal{C})$, then for every playable coalition strategy $\nu \in \mathcal{S}_J$ of $\mathcal{C}$ it holds that (J, ν) is a possible deal.

Since every safely playable strategy of a confirmed deal $(J, \mathcal{C}')$ is in particular a playable strategy of the stable set $\mathcal{C}'$, we immediately get the following result.

Corollary 12
Let $(J, \mathcal{C}')$ be a confirmed deal based on a set of commitments $\mathcal{C}$, then for every safely playable coalition strategy $\nu \in \mathcal{S}_J$ of $(J, \mathcal{C}')$ it holds that (J, ν) is a possible deal.

4.2 Nash Equilibria

To illustrate how the introduced notions can be applied, we now show that agents can choose their commitments such that the union of the deals is a PNE (if one exists). The following results link the best response of a debtor to the best responses of the other agents involved in a deal.

Proposition 13
Let $\mathcal{C}$ be a set of commitments in the BG G such that every $c \in \mathcal{C}$ respects the debtor's best responses and let (J, ν) be a possible deal based on $\mathcal{C}$. For every $i \in J$ and every $\nu' \in \mathcal{S}_{N \setminus J}$ it holds that $\mathcal{BR}(G, \nu_i, (\nu_{-i}, \nu'))$.

Intuitively, Proposition 13 expresses that every possible deal based on commitments which respect the debtor's best responses has the property that, regardless of what the agents who are not part of the deal decide to do, the agents who are part of the deal play a best response by playing the strategy specified in the deal.

Corollary 14
Let $\mathcal{C}$ be a set of commitments in the BG G such that every $c \in \mathcal{C}$ respects the debtor's best responses and let $(J, \mathcal{C}')$ be a confirmed deal based on $\mathcal{C}$. For every $i \in J$, every safely playable coalition strategy $\nu \in \mathcal{S}_J$ and every $\nu' \in \mathcal{S}_{N \setminus J}$ it holds that $\mathcal{BR}(G, \nu_i, (\nu_{-i}, \nu'))$.

We can now straightforwardly derive the following proposition from Proposition 4 and Corollaries 6 and 14.

Proposition 15
Suppose that agents only announce consistent sets of commitments respecting their best responses. If all agents are part of a confirmed deal, in either the original BG or one of its reductions based on previously closed deals, then the union of the safely playable coalition strategies is a PNE.

Note that, by definition, no agent has an incentive to individually deviate from an obtained deal which is a PNE. We moreover have the following result.

Proposition 16
If the BG G has a PNE, then there exists a sequence of commitments such that every agent is guaranteed to become part of a confirmed deal and the union of these deals is a PNE, without requiring prior knowledge of the other agents' goals.

Note, however, that the existence of a PNE is not required to obtain a confirmed deal. Consider for instance a 2 agent BG with $\Phi_i = \{d_i\}$, $\delta_i = \top$, $\Gamma_1 = \{d_1 \wedge d_2; \neg d_1 \wedge \neg d_2\}$ and $\Gamma_2 = \{\neg d_1 \wedge \neg d_2; d_1 \wedge d_2\}$. In other words, agent 1 prefers disarmament of both nations over arming both nations, and vice versa for agent 2. If the agents announce the commitments $c_1 = (1; 2; d_2; d_1)$ and $c_2 = (2; 1; \neg d_1; \neg d_2)$ in the first round, there are no possible deals. If they additionally announce the commitments $c_1' = (1; 2; \neg d_2; \neg d_1)$ and $c_2' = (2; 1; d_1; d_2)$ in the second round, 2 confirmed deals involving both agents are obtained, namely $(\{1,2\}, \{c_1', c_2\})$ and $(\{1,2\}, \{c_1, c_2'\})$. The unique safely playable strategies are respectively $\{\neg d_1, \neg d_2\}$, i.e. both nations arm, and $\{d_1, d_2\}$, i.e. both nations disarm. These two outcomes are Pareto optimal, i.e. no other outcome exists such that both agents would be better off.

4.3 Commitment-based Deals in Negotiation Protocols

The formalized concepts of commitments and commitment-based deals can easily be plugged into existing negotiation protocols. This allows to select protocols that best suit the needs of the application. As an illustration, we configure two kinds of protocols.

Firstly, we explain how our notions can be used in a multilateral monotonic concession protocol [34, 18], in which agents incrementally make concessions to reach an agreement. Given our framework of BGs with prioritized goal bases, this is a very intuitive approach for the agents: if a commitment corresponding to the first priority goal of an agent does not lead to any deal, the agent can concede by considering its second priority goal as well.

Example 10
Reconsider the context of disarmament of nations and suppose a pacifistic agent 1 is in a BG with 3 agents and has the goal base $\Gamma_1 = \{d_1 \wedge d_2 \wedge d_3; d_1 \wedge (d_2 \vee d_3); d_1\}$. Its first commitment would be $c_1 = (1; \{2,3\}; d_2 \wedge d_3; d_1)$, coinciding with its first priority goal of global disarmament. Now assume that agents 2 and 3 have communicated commitments such that not a single possible deal arose, e.g. $c_2 = (2; \{1,3\}; \neg d_1 \vee \neg d_3; \neg d_2)$ and $c_3 = (3; 1; d_1; d_3)$. Then agent 1 can concede by communicating the commitment $c_1' = (1; \{2,3\}; d_2 \vee d_3; d_1)$, coinciding with its second priority goal.

Concessions can thus easily be captured by opening with commitments corresponding to the first priority goal, then conceding to the disjunction of the first and second priority goal, next to the disjunction of the first, second and third priority goal etc. Note that in Example 10, the disjunction of agent 1's first and second priority goal is equivalent with its second priority goal.
The notion of a confirmed deal can fulfill the concept of agreement used in the monotonic concession protocol [34, 18]. In Example 10, the concession of agent 1 would lead to the confirmed deal $(\{1,3\}, \{c_1', c_3\})$ based on $\{c_1', c_2, c_3\}$, even without concession of agents 2 and 3. The corresponding unique safely playable coalition strategy is $\{d_1, d_3\}$.
Recall that e.g. proposals corresponding to propositional formulas would not be sufficient to obtain deals between coalitions, as they do not take the control assignment of the action variables into account. Finally note that previous work on monotonic commitment concession does not address goal-related concession [42].

As another example, consider the negotiation protocol described in [33], in which a broker agent matches proposals, and then notifies the agents — which submitted the proposals — about the possibility of agreement. It is easy to see that, by using commitments as proposals, we can straightforwardly use our definitions of a deal to capture the notion of matching proposals. In [33], the agents are supposed to negotiate about which of the possible agreements is to be chosen after they received the notification. However, the presence of a broker agent — a central entity — can also bypass the need of this extra negotiation by selecting at most one deal per agent. It could even use a social criterion, e.g. by selecting the deals which involve the largest number of agents. Whether such interventions are desirable or even justifiable strongly depends on the context of the application.

5 Related Work

To the best of our knowledge, this paper is the first to study commitments in BGs. Therefore, we structure this section by addressing the related work w.r.t. BGs and commitments separately.

5.1 Related Work w.r.t. Boolean Games

Our study on opportunities for agreement and the formation of coalitions is reminiscent of cooperative game theory [2]. The study of Boolean games from a cooperative point of view has led to a variety of concepts, such as e.g. the core, stable sets of outcomes [17], efficient coalitions, weak and strong core [8] and stable coalitions [35]. In this existing work only BGs with a single goal are considered, yet costs associated with strategies are used to obtain non-binary utilities. In this paper, however, the coalition concepts, i.e. the deals, are based on a set of commitments instead of on the BG. Nonetheless, since it is sensible that agents' commitments are related to their goals (see Section 3), it is likely that, under certain assumptions, links can be found between the different concepts. For instance, in [8] a coalition is called efficient iff it can satisfy the goal of every coalition partner, regardless of the strategies of the agents outside the coalition. It is clear that our concepts offer a way to obtain the efficient coalitions in the case of BGs with a single goal: if every agent communicates commitments coinciding with their goal, the efficient coalitions are exactly the possible deals. Analogously, negotiation strategies might be characterized such that agents obtain an agreement corresponding to alternative solution concepts, such as the weak and strong core [8], as these can be linked to the concept of efficient coalitions. Further investigation of these links, however, lies beyond the scope of this paper.

Although negotiation [29, 1] and BGs [9, 17, 14, 7, 8, 41, 5] have been widely studied, only few works have considered protocols in BGs which allow agents to actually coordinate towards agreeable outcomes. A multilateral negotiation protocol for BGs has been investigated in [17], where it is shown that, when the logical structure of the goals is restricted to positive goals (i.e. no connectives other than $\wedge$ and $\vee$ are used), the protocol is guaranteed to end in a Pareto optimal outcome, meaning that no agent can improve its position without another agent being worse off. BGs have also been extended to endogenous variants [39], involving a pre-play negotiation phase, in which agents can try to influence the decisions of other agents by means of side payments, i.e. transferring parts of their utility to other agents. In [19], an extension of BGs is used to model voting strategies in binary aggregation. A multilateral negotiation protocol for BGs with prioritized goals has been developed in [15]. The protocol is driven by an agreement rule which guarantees a fair and efficient outcome under complete knowledge about the other agents' goals. It is extended to BG settings with incomplete knowledge, in which case a weaker notion of fairness and efficiency is guaranteed. In this agreement protocol, the order of the agents strongly influences the outcome. In contrast, in this paper we introduce building blocks for multilateral negotiation protocols in which arbitrary propositional formulas can be used to specify goals. Moreover, the building blocks can be used in protocols in which the ordering of the agents does not influence the agreements, such as the multilateral protocol with the broker agent. Furthermore, it is easy to see that the assumption of transferable utility or knowledge about the other agents' goals is not required to obtain deals.

5.2 Related Work w.r.t. Commitments

Outside the setting of BGs, the characterization of commitments and commitment-based protocols has been extensively studied [36, 25, 40, 16, 31, 32, 37, 12, 28, 22, 38]. Commitments, which can be considered a form of pre-play moves [25], are studied from a game-theoretical perspective in Schelling's seminal work [36]. Schelling explains how commitments can be threats, when one commits to deviate from its own best response to damage the opponent, or promises, when one commits to deviate from its own best response to cooperate with the opponent [36]. However, since we assume no knowledge about the other agents' goals in this paper, the agents do not have the required information to determine whether their commitments are either one of these. Instead, commitments here should be understood primarily as a way to communicate its own incentives, a point of view which is also noted by Schelling [36].

In this paper, we have introduced the concept of a consistent set of commitments. A series of postulates for commitments has been described in [37], in which the debtor and creditor are fixed. Using the *monotonicity* postulate, it is clear that an inconsistent set of commitments with a fixed debtor and creditor would violate the *strong consistency* postulate. To obtain the notion of consistency as defined in our paper though, a meaningful generalization of the postulates to variable creditors is required. In [22], conflicts between commitments are discussed, where it is assumed that the domain dependent conflict knowledge is already present. For instance, the same car cannot be rented simultaneously by two different individuals. The framework in [22] is more involved than ours, as it relies on event calculus to formalize conflicts between commitments. The debtor and creditor are irrelevant in the detection of conflicts [22]. Three different notions of conflict are defined, but even if we fix the debtor, none these notions coincide with our concept of consistency, as none of them seem to take into account whether the antecedents can simultaneously be satisfied or not. In [23], the notion of feasibility of commitments is introduced. The intuitive idea behind this concept is the same as ours w.r.t consistency, i.e. checking whether it is possible for an agent to fulfill a set of commitments all together [23]. Their elaboration, however, is slightly different from ours: the feasibility of the commitments of an agent i does not only take the commitments with debtor i into account, but also those with creditor i. For instance, if you have committed to make two payments to two different agents but only have sufficient money for one of the payments, we consider this pair of commitments inconsistent. In [23], this set might still be considered feasible in case there is e.g. a commitment of a reliable agent to make a payment to you. Moreover, for a set of commitments to be feasible there should exist an execution that discharges them all, i.e. that brings about all the consequences [23]. Translated to our framework, that would imply that a pair of commitments of the form $(1; 2; d_2; d_1)$ and $(1; 2; \neg d_2; \neg d_1)$ is not feasible. In this paper however, this pair is consistent, due to the inconsistency of the antecedents.

Several studies on the relationship between goals and commitments can be found in the multi-agent literature [11, 38, 20]. The variety among this work can be partially explained by the usage of different representations for the goals. The goal models in [11] are specified in Tropos, an agent-oriented software engineering framework [10]. Goals can either be decomposed as conjunctions or as disjunctions of subgoals. Moreover, the link between the possible achievement of pairs of goals is formalized. In this framework, goals are not prioritized. The authors exploit commitments for goal support and vice versa and provide elaborated semantics. In [38], goals consists of a precondition, an in-condition, a post-condition, a success condition and a failure condition. An agent can have multiple goals, but it is assumed that they are mutually consistent. No priority between goals is used. A formalization from goal to commitment and vice versa is provided, based on guarded rules. These rules operate on the goals (e.g. consider, activate, suspend, reactivate, drop) and commitments (e.g. release, cancel). In [20], goals are represented in a similar way as in [38], but the in-condition and post-condition are dropped. The framework also involves beliefs about the other agents' goals and capabilities. It is likely that BGs could be embedded into this framework in case it were to be generalized to allow priorities in the set of goals. As in [11], a notion of goal support is introduced and the authors provide algorithms to generate commitments to support their goals [20, 27, 21]. To this end, agents also use their beliefs about the goals of other agents to formulate commitments which are more likely to be accepted [20]. The notion of goal support takes all commitments into account and considers a goal to be supported if there is a chain of commitments leading to the satisfaction of the goal. For example, if agent 1 and 2 have communicated the commitments $(1; 2; d_2; d_1)$ and $(2; 1; d_1; d_2)$, then none of these agents is considered to support the goal $\gamma = d_1 \wedge d_2$. In contrast, we consider these as commitments coinciding with γ, which moreover form a stable set to bring about γ. Investigating the possible formal links between the aforementioned work and ours is an interesting topic for further research.

A large amount of the prior work on commitment-based protocols mainly focuses on practical aspects, e.g. investigating the life-cycle of commitments [40], solving misalignment [12], detecting exceptions in commitments [28], applications in a business context [16], etc., and pays less attention to the underlying multi-agent system. In contrast, we analyze commitment-based deals specifically in the context of BGs, allowing us to exploit game-theoretical concepts such as utility and best response to define sensible commitments and to introduce suitable notions of deals between coalitions of agents.

6 Conclusion

The aim of this paper was to study, at a general level, the notion of commitments and commitment-based deals in Boolean games. First we have formalized the notion of commitments and explained how goals can be related to sets of commitments. Then, we investigated commitment-based deals, which rely on the idea of identifying stable sets among the commitments made by a group of agents. Finally, we have illustrated how the notions of commitments and deals can be used to guarantee e.g. the Nash property in the obtained deals, or to configure existing negotiation protocols. The latter allows a flexible use of the introduced concepts: depending on the application context one can, for instance, either plug our building blocks into a distributed or a centralized protocol.

Note that in a practical implementation, additional aspects might need to be addressed. However, to a large extent we can rely on the available results w.r.t. (algebraic) formalization of commitments [32, 37, 38] and commitment-based protocols [40, 31, 12]. For instance, to escape deadlock agreements (characterized by cyclic dependencies) or avoid the possibility of cheating caused by imperfect synchrony, one could respectively rely on the 2PC protocol [40] and the lockstep-protocol [4]. Deciding which confirmed deal is closed when more than one arises should also be addressed (e.g. choosing the largest coalition).

Finally, an important issue is how agents can act strategically in how they formulate commitments. This should be based on their beliefs about the other agents' goals, their risk aversion and/or their readiness to concede, and could be implemented using techniques such as Monte Carlo tree search. Alternatively, an approach similar to the one in [20] can be investigated, where the beliefs about the agents' goals are used to generate commitments that are more likely to be accepted by other agents.

REFERENCES

[1] L. Amgoud and H. Prade, 'A possibilistic logic modeling of autonomous agents negotiation', in *Progress in Artificial Intelligence*, 360–365, Springer, (2003).

[2] R.J. Aumann, 'Game theory', in *The New Palgrave*, eds., J. Eatwell, M. Milgate, and P. Newman, 1–53, Macmillan, London and Basingstoke, (1989).

[3] C. Bartolini, C. Preist, and N.R. Jennings, 'A software framework for automated negotiation', in *Software Engineering for Multi-Agent Systems III*, 213–235, Springer, (2004).

[4] N.E. Baughman and B.N. Levine, 'Cheat-proof playout for centralized and distributed online games', in *INFOCOM 2001. Twentieth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, volume 1, pp. 104–113, (2001).

[5] J. Ben-Naim and E. Lorini, 'Evaluating power of agents from dependence relations in Boolean games', in *Proc. AAMAS*, pp. 853–860, (2014).

[6] E. Bonzon, M.-C. Lagasquie-Schiex, and J. Lang, 'Compact preference representation for Boolean games', in *Proc. PRICAI*, pp. 41–50. Springer-Verlag, (2006).

[7] E. Bonzon, M.-C. Lagasquie-Schiex, and J. Lang, 'Dependencies between players in Boolean games', in *Proc. ECSQARU*, volume 4724 of *LNCS*, pp. 743–754. Springer, (2007).

[8] E. Bonzon, M.-C. Lagasquie-Schiex, and J. Lang, 'Effectivity functions and efficient coalitions in Boolean games', *Synthese*, **187**, 73–103, (2012).

[9] E. Bonzon, M.-C. Lagasquie-Schiex, J. Lang, and B. Zanuttini, 'Boolean games revisited', in *Proc. ECAI*, pp. 265–269. ACM, (2006).

[10] P. Bresciani, A. Perini, P. Giorgini, F. Giunchiglia, and J. Mylopoulos, 'Tropos: An agent-oriented software development methodology', *Autonomous Agents and Multi-Agent Systems*, **8**(3), 203–236, (2004).

[11] A.K. Chopra, F. Dalpiaz, P. Giorgini, and J. Mylopoulos, 'Reasoning about agents and protocols via goals and commitments', in *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*, volume 1, pp. 457–464. International Foundation for Autonomous Agents and Multiagent Systems, (2010).

[12] A.K. Chopra and M.P. Singh, 'Multiagent commitment alignment', in *Proceedings of AAMAS*, pp. 937–944. International Foundation for Autonomous Agents and Multiagent Systems, (2009).

[13] A. Darwiche, 'Decomposable negation normal form', *Journal of the ACM (JACM)*, **48**(4), 608–647, (2001).

[14] S. De Clercq, S. Schockaert, M. De Cock, and A. Nowé, 'Possibilistic Boolean games: Strategic reasoning under incomplete information', in *Proc. JELIA*, pp. 196–209, (2014).

[15] S. De Clercq, S. Schockaert, A. Nowé, and M. De Cock, 'Multilateral negotiation in Boolean games with incomplete information using generalized possibilistic logic', in *Proceedings of IJCAI*, pp. 2890–2896, (2015).

[16] N. Desai, A.U. Mallya, A.K. Chopra, and M.P. Singh, 'Interaction protocols as design abstractions for business processes', *Software Engineering, IEEE Transactions on*, **31**(12), 1015–1027, (2005).

[17] P.E. Dunne, W. van der Hoek, S. Kraus, and M. Wooldridge, 'Cooperative Boolean games', in *Proc. AAMAS*, volume 2, pp. 1015–1022. IFAAMAS, (2008).

[18] U. Endriss, 'Monotonic concession protocols for multilateral negotiation', in *Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*, pp. 392–399. ACM, (2006).

[19] U. Grandi, D Grossi, and P. Turrini, 'Equilibrium refinement through negotiation in binary voting', in *Proceedings of IJCAI*, pp. 540–546, (2015).

[20] A. Günay, M. Winikoff, and P. Yolum, 'Commitment protocol generation', in *Declarative agent languages and technologies X*, 136–152, Springer, (2012).

[21] A. Günay, M. Winikoff, and P. Yolum, 'Dynamically generated commitment protocols in open systems', *Autonomous Agents and Multi-Agent Systems*, **29**(2), 192–229, (2015).

[22] A. Günay and P. Yolum, 'Detecting conflicts in commitments', in *Declarative Agent Languages and Technologies IX*, 51–66, Springer, (2011).

[23] A. Günay and P. Yolum, 'Constraint satisfaction as a tool for modeling and checking feasibility of multiagent commitments', *Applied intelligence*, **39**(3), 489–509, (2013).

[24] P. Harrenstein, W. van der Hoek, J.-J. Meyer, and C. Witteveen, 'Boolean games', in *Proc. TARK*, pp. 287–298. MKP Inc., (2001).

[25] J. Hirshleifer, 'Game-theoretic interpretations of commitment', *Evolution and the capacity for commitment*, 77–93, (2001).

[26] N.R. Jennings, P. Faratin, A.R. Lomuscio, S. Parsons, M.J. Wooldridge, and c. Sierra, 'Automated negotiation: prospects, methods and challenges', *Group Decision and Negotiation*, **10**(2), 199–215, (2001).

[27] O. Kafalı, A. Günay, and P. Yolum, 'Gosu: Computing goal support with commitments in multiagent systems', in *Proceedings of 21st European Conference on Artificial Intelligence*, pp. 477–482, (2014).

[28] Ö. Kafalı and P. Yolum, 'Detecting exceptions in commitment protocols: Discovering hidden states', in *Languages, Methodologies, and Development Tools for Multi-Agent Systems*, 112–127, Springer, (2009).

[29] S. Kraus, 'Negotiation and cooperation in multi-agent environments', *Artificial intelligence*, **94**(1), 79–97, (1997).

[30] J. Lang, P. Liberatore, and P. Marquis, 'Propositional independence-formula-variable independence and forgetting', *Journal of Artificial Intelligence Research*, 391–443, (2003).

[31] A.U. Mallya and M.P. Singh, 'Introducing preferences into commitment protocols', in *Agent Communication II*, 136–149, Springer, (2006).

[32] A.U. Mallya and M.P. Singh, 'An algebra for commitment protocols', *Autonomous Agents and Multi-Agent Systems*, **14**(2), 143–163, (2007).

[33] Piotr Pałka, 'Multilateral negotiations in distributed, multi-agent environment', in *Computational Collective Intelligence. Technologies and Applications*, 80–89, Springer, (2011).

[34] Jeffrey S Rosenschein and Gilad Zlotkin, *Rules of encounter: designing conventions for automated negotiation among computers*, MIT press, 1994.

[35] L. Sauro and S. Villata, 'Dependency in cooperative Boolean games', *Journal of Logic and Computation*, **23**(2), 425–444, (2013).

[36] T.C. Schelling, *The strategy of conflict*, Harvard university press, 1980.

[37] M.P. Singh, 'Semantical considerations on dialectical and practical commitments.', in *AAAI*, volume 8, pp. 176–181, (2008).

[38] P.R. Telang, M.P. Singh, and N. Yorke-Smith, 'Relating goal and commitment semantics', in *Programming Multi-Agent Systems*, 22–37, Springer, (2012).

[39] P. Turrini, 'Endogenous games with goals: side-payments among goal-directed agents', *Autonomous Agents and Multi-Agent Systems*, 1–28, (2015).

[40] F. Wan and M.P. Singh, 'Formalizing and achieving multiparty agreements via commitments', in *Proceedings of AAMAS*, pp. 770–777. ACM, (2005).

[41] M. Wooldridge, U. Endriss, S. Kraus, and J. Lang, 'Incentive engineering for Boolean games', *Artif. Intell.*, **195**, 418–439, (2013).

[42] P. Yolum and M.P. Singh, 'Enacting protocols by commitment concession', in *Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems*, pp. 116–123. ACM, (2007).