

Boolean Negotiation Games

Nils Bulling¹ and Koen V. Hindriks¹

Abstract. We propose a new strategic model of negotiation, called Boolean negotiation games. Our model is inspired by Boolean games and the alternating offers model of bargaining. It offers a computationally grounded model for studying properties of negotiation protocols in a qualitative setting. Boolean negotiation games can yield agreements that are more beneficial than stable solutions (Nash equilibria) of the underlying Boolean game.

1 Introduction

There are at least two prominent methodologies to analyse negotiations between agents: off-line, e.g. [8], using game theoretic techniques, and online, e.g. [2], using heuristic and evolutionary models. Most work in the game theoretic approach is based on Rubinstein's bargaining model of alternating offers [9], often making the assumptions of perfect rationality and perfect information.

We propose a new, compact model that allows us to investigate strategic aspects of negotiation protocols. The model we propose is inspired by Boolean games [4] (BG) which have become a popular model in the multi-agent domain. The many variants and extensions of BGs related to, e.g., knowledge [1], control and manipulation [6, 7], secret goals [5], dependencies [3], and pre-play negotiations about payoffs [10], just to name a few, make them an ideal starting point for our purposes. Boolean games, for example, allow us to also study aspects of control and power in a negotiation.

The main contributions of the work we present here consist of a model of negotiation called *Boolean negotiation game* (BNG) and a formal analysis of a protocol that does not allow repeating offers using this model. In the context of BGs the non-repetition of offers naturally yields finite games, which arise in many practical contexts. We introduce negotiation equilibria and are able to show that they always exist and that they can yield agreements which are more beneficial than the Nash equilibria of the underlying BG. In this context, the negotiation protocol plays a crucial role. A negotiation protocol gives rise to a specific unfolding of a BG with similarities to extensive games, but this unfolding is more general as it may not result in a complete agreement on all outcomes resulting in a smaller BG being played after the negotiation phase. As such, different properties of negotiation protocols greatly affect the game being played.

2 Boolean Negotiation Games

Boolean negotiation games (BNGs) allow players to interact in BGs by exchanging proposals sequentially. A negotiation protocol is imposed on a given BG affecting the possible actions of agents. Such a protocol adds a new layer of strategic interaction as not just the plain selection of a specific proposal is important but also the timing

is of crucial significance, for example, in the setting where proposals cannot be offered more than once.

In order to introduce our model we first define the notion of a *generalized extensive Boolean game* (GBG). These games generalise the standard game theoretic notion of an extensive (Boolean) game [9]: (i) agents' actions are not limited to setting a single variable at a time but may *propose settings for multiple variables*, in principle including those which they *do not control*, (ii) at terminal histories, not all variables must be assigned a truth assignment. We assume a vector **Act** consisting of sets of actions. The idea is that agent i draws its actions from the set **Act** _{i} , where the agent's valuations of propositions are typically also taken to be actions. We say that an action act_1 is said to *conflict* with action act_2 performed earlier during the game if the actions assign different truth values to a proposition. As in extensive form games, a protocol determines which agent's turn it is as well as the enabled actions at the current situation. A (possibly empty) sequence of actions is called a *history*. A *protocol* P maps a history h to a tuple (i, V) indicating that it is agent i 's turn in h and that actions $V \subseteq \text{Act}_i$ are enabled.

A GBG is a BG together with a protocol P . Given a protocol P , a P -strategy for agent i for that game is a (partial) function π_i the domain of which consists of all non-terminal P -histories h — histories consistent with P — at which it is player i 's turn and which assigns a P -enabled action to such histories. A *profile* of P -strategies π yields a unique P -run ρ_π , i.e. a sequence of actions consistent with the strategy.

The unique P -run ρ_π yielded by strategy profile π may not set the truth of all variables. In addition to that, some of the performed actions might be conflicting, e.g. a player may set a variable true and the same variable false later during the game. To determine an outcome of a GBG, we need to resolve these conflicts. The general rule that we use to determine the outcome is that any action that conflicts with an action performed later during the game is reverted and ignored in the computation of the outcome. Moreover, if the resulting outcome ξ is not a (full) valuation for all variables in Π , then the agents need to settle on the remaining variables in some other way. To settle on variables for which the agents did not settle on a valuation, the agents establish values by means of the ξ -reduced BG of the GBG, which is the BG in which each variable the truth value of which is defined by ξ is replaced by that truth value. Thus, a strategy of an agent in a GBG also needs to define which variables the agent sets in the resulting reduced BG. A formal treatment is out of the scope of this abstract, however, once this is defined formally it allows to introduce the notion of a *generalized equilibrium* taking into account that a strategy consists of two parts: a P -strategy and a strategy for the reduced BGs. We observe that a generalised equilibrium does not have to exist as some reduced BGs may not have a stable outcome. To see this, consider the trivial case in which the GBG consists of a single root node after which normal form games not having any Nash

¹ TU Delft, The Netherlands, email: {n.bulling, k.v.hindriks}@tudelft.nl.

equilibria are played: then, there is no generalised equilibrium either.

A *Boolean negotiation game* (BNG) is a special GBG in which the underlying protocol satisfies specific properties tailored towards negotiation settings. It is well accepted that there are at least two minimal requirements most negotiations should satisfy: (i) agents can make proposals and are able to respond to them; (ii) agents need to approve a possible agreement before it is concluded [9]. The latter point also implies that taking part in a negotiation should be individually rational for each agent. Based on these two properties we now introduce negotiation protocols and BNGs. In the negotiation setting, actions should be thought of as *proposals* made to the other players. As a consequence, a proposal conflicting with a proposal made earlier implicitly *rejects* the earlier proposal and serves the purpose of a *counter-proposal*. To implement point (ii) above, we identify a sub-class of protocols that requires all agents except for the agent who made the last proposal to explicitly approve the agreement that is on the table. Therefore, players which are happy with the current proposals can accept. Players have to be cautious, though, because if a proposal is made all other players could accept it which concludes the negotiation. To this end, the empty valuation ξ_\emptyset is now interpreted as an accept action and is identified with the special action accept. In order to allow agents to accept, the protocol needs to support this. In that case we say that a protocol *supports agreeing*. Similarly a protocol *supports quitting* if a quit action allowing the agent to leave the negotiation is enabled after each non-terminal history. Then, a run is *closing* iff its final action is quit and it does not contain any further quit actions. There are different approaches how to deal with a quitting agent. We assume that in that case no deal is reached and the whole negotiation ends. Finally, a *negotiation protocol* (NP) is a protocol P which is turn-taking (i.e. players act in turns), supports agreeing as well as quitting and in which each P -run ρ is either agreeing containing no quit action, or is closing. Consequently, a BNG is a GBG including a negotiation protocol.

The definition of NPs is very general. It often makes sense to put some restrictions on the proposals that can be made. For example, it is usually not helpful to make the same proposal again and again. If a proposal has not resulted in an agreement it will, under reasonable assumptions, also not do so if it is made over again, only if the proposer counts on wearing out his/her opposite. The assumption is also reasonable in terms of real negotiations where it is often difficult to get back to a previously rejected proposal. Therefore, we focus on non-repeating protocols. A NP P is *non-repeating* if no proposal can be made twice with the exception of ξ_\emptyset playing the role of the accept action. Furthermore, in order to investigate agents' interactions we focus on two types of protocols: one in which agents make proposals concerning their own variables only; and one where agents propose full valuations only. In the following let \mathcal{N} be a BNG.

We are especially interested in the question whether agents have an \mathcal{N} -strategy profile σ which yields an agreement which is acceptable for all agents, given the possible outcomes of the underlying BG of \mathcal{N} . As usual in negotiation settings agents have a reservation value which corresponds to the payoff below which a player would refuse any proposal. A rather strict notion of reservation value would be a player's maxmin-strategy defining an outcome which the player can guarantee on its own. We call the corresponding reservation value the *maxmin reservation value*. In the strategic setting we consider here, it makes good sense to relate the reservation value to outcomes of Nash equilibria, as they give a payoff at least as good as the maxmin reservation value. In general, there can be more than one Nash equilibrium, therefore, we define a weak and a strong notion. The *greedy reservation value* (resp. *modest reservation value*) is the

player's maximal (resp. minimal) payoff received by any Nash equilibrium. If a game does not have any Nash equilibria both reservation values are defined as the player's maxmin reservation value. We refer to *greedy* (resp. *modest*) agents as such which use as baseline their greedy (resp. modest) reservation values. Whereas the existence of a subgame perfect equilibrium in finite extensive form games is guaranteed by Kuhn's theorem (cf. [9]) this is not obvious in our setting. Indeed, it does not hold for the notion of generalised equilibrium put forward in the context of GBGs. The reason is that agents can quit the negotiation which results in a Boolean game over the not yet fixed variables. This BG may not have any Nash equilibria which also explains why a generalised equilibrium may not exist. In general the solution concept of generalised equilibrium is too strong as players base their decision on reservation values. Therefore, we introduce the weaker solution concept of a *negotiation equilibrium*. This solution concept makes no further assumption about the players' behavior if a (complete) agreement is not reached apart from assuming that each player can be ensured to receive a payoff at least as good as its reservation value in the resulting reduced BG. We can show that such an equilibrium always exists.

3 Future Work

We have proposed a formal framework for studying strategic aspects of negotiations in the compact framework of Boolean games. We used this to study negotiation protocols that do not allow repeating offers. There are many other interesting constraints on protocols that we could study within our framework. There are also many questions that we would like to study in more detail. For example, which protocols guarantee Pareto optimal outcomes? We are also particularly interested in studying negotiation with partial knowledge. Although a lot of research on negotiation with incomplete information has already been done, analysing this setting theoretically remains a challenge. Our model provides a starting point for creating a theoretical model of negotiation with incomplete information in future work.

REFERENCES

- [1] Thomas Ågotnes, Paul Harrenstein, Wiebe van der Hoek, and Michael Wooldridge, 'Boolean games with epistemic goals', in *Logic, Rationality, and Interaction*, 1–14, Springer, (2013).
- [2] Tim Baarslag, Mark J.C. Hendriks, Koen V. Hindriks, and Catholijn M. Jonker, 'Learning about the opponent in automated bilateral negotiation: a comprehensive survey of opponent modeling techniques', *Autonomous Agents and Multi-Agent Systems*, 1–50, (2015).
- [3] Elise Bonzon, Marie-Christine Lagasque-Schiex, and Jérôme Lang, 'Dependencies between players in boolean games', *International Journal of Approximate Reasoning*, **50**(6), 899–914, (2009).
- [4] Elise Bonzon, Marie-Christine Lagasque-Schiex, Jérôme Lang, and Bruno Zanuttini, 'Boolean games revisited', in *ECAI*, pp. 265–269, (2006).
- [5] Nils Bulling, Sujata Ghosh, and Rineke Verbrugge, 'Reaching your goals without spilling the beans: Boolean secrecy games', in *PRIMA*, pp. 37–53, Dunedin, New Zealand, (2013).
- [6] Ulle Endriss, Sarit Kraus, Jérôme Lang, and Michael Wooldridge, 'Designing incentives for boolean games', in *AAMAS*, pp. 79–86, (2011).
- [7] John Grant, Sarit Kraus, Michael Wooldridge, and Inon Zuckerman, 'Manipulating boolean games through communication', in *IJCAI*, pp. 210–215, (2011).
- [8] Sarit Kraus, *Strategic Negotiation in Multiagent Environments*, MIT Press, 2001.
- [9] Martin J. Osborne and Ariel Rubinstein, *A Course in Game Theory*, The MIT Press, 1994.
- [10] Paolo Turrini, 'Endogenous boolean games', in *Proceedings of the Twenty-Third international joint conference on Artificial Intelligence*, pp. 390–396. AAAI Press, (2013).