Searching for Pure Strategy Equilibria in Bilateral Bargaining With One-sided Uncertainty

(Extended Abstract)

Bo An
Dept. of Computer Science
University of Massachusetts, Amherst, USA
ban@cs.umass.edu

Nicola Gatti
Dip. Elettronica e Inforazione
Politecnico di Milano, Milano, Italy
ngatti@elet.polimi.it

Victor Lesser
Dept. of Computer Science
University of Massachusetts, Amherst, USA
lesser@cs.umass.edu

ABSTRACT

The problem of finding agents' rational strategies in bargaining with incomplete information is well known to be challenging. The literature provides a collection of results for very narrow uncertainty settings, but no generally applicable algorithm. In this paper, we focus on the alternating-offers finite horizon bargaining protocol with one-sided uncertainty regarding agents' reserve prices. We provide an algorithm based on the combination of game theoretic analysis and search techniques which finds agents' equilibrium in pure strategies when they exist. Our approach is sound, complete and, in principle, can be applied to other uncertainty settings.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multiagent Systems

General Terms

Algorithms, Economics

Keywords

Negotiation, sequential equilibrium, uncertainty

1. INTRODUCTION

This paper focuses on finding agents' rational strategies in incomplete information bilateral bargaining. We consider the most common bargaining protocol, i.e., the Rubinstein's alternating-offers [5], which has been widely used in the bargaining theory literature, e.g., [3, 6]. We analyze the situation with one-sided uncertain reserve prices and where agents have deadlines. The microeconomic literature provides a number of closed form results with very narrow uncertainty settings. For instance, Rubinstein [6] considered bilateral infinite horizon bargaining with uncertainty over two possible discount factors. Gatti *et al*. [3] analyzed bilateral bargaining with one-sided uncertain deadlines. Chatterjee and Samuelson [2] studied bilateral *infinite* horizon bargaining with two-type uncertainty over the reservation values. The absence of agents' deadlines makes these two results nonapplicable to the situation we study in the paper. An *et al*. [1] only consider two-type uncertainty about reserve prices. The presence of many types increases the computational complexity of the procedure to find equilibrium strategies and requires more stringent equilibrium existence conditions. Operations research inspired algorithms such as Miltersen-Sorensen [4] work only on games with finite number of strategies, and therefore cannot be applied to bargaining in which each agent's strategy space is continuous.

We develop a novel algorithm to find a pure strategy sequential equilibrium in bilateral bargaining with multi-type uncertainty. Our algorithm combines together game theoretic analysis with state space search techniques and it is sound and complete. Our approach is based on the following two observations: 1) with pure strategies, agents' possible choice rules regarding whether different buyer types will behave in the same way or in different ways at a decision making point are finite, and 2) given a tree of choice rules (each time point is assigned a choice rule) we are able to derive theoretically the agents' optimal strategies (by a Bayesian extension of backward induction) and to check whether or not a sequential equilibrium there is with such tree of belief systems.

2. ONE-SIDED UNCERTAINTY

We consider the discrete time bargaining between a buyer $\mathbf{b}$ and a seller $\mathbf{s}$. The seller wants to sell a single indivisible good to the buyer. All the agents enter the market at time 0. An alternating-offers bargaining protocol is utilized. Formally, the buyer $\mathbf{b}$ and the seller $\mathbf{s}$ can act at times $t \in \mathbb{N}$. The player function $\iota : \mathbb{N} \rightarrow \{\mathbf{b}, \mathbf{s}\}$ returns the agent that acts at time t and is such that $\iota(t) \neq \iota(t+1)$, i.e., a pair of agents bargain by making offers in alternate fashion.

Possible actions $\sigma^t_{\iota(t)}$ of agent $\iota(t)$ at any time point $t > 0$ are: 1) $\mathit{offer}[x]$, where $x \in \mathbb{R}$ is the proposed price; 2) exit, which indicates that negotiation fails; and 3) accept, which indicates that $\mathbf{b}$ and $\mathbf{s}$ have reached an agreement. At time point $t = 0$ the only allowed actions are 1) and 2). If $\sigma^t_{\iota(t)} = \mathit{accept}$ the bargaining stops and the outcome is (x, t), where x is the value such that $\sigma^{t-1}_{\iota(t-1)} = \mathit{offer}[x]$. This is to say that the agents agree on the value x at time point t. If $\sigma^t_{\iota(t)} = \mathit{exit}$ the bargaining stops and the outcome is $FAIL$. Otherwise the bargaining continues to the next time point.

Each agent $\mathbf{a} \in \{\mathbf{b}, \mathbf{s}\}$ has a utility function $U_{\mathbf{a}} : (\mathbb{R} \times \mathbb{N}) \cup FAIL \rightarrow \mathbb{R}$, which depends on $\mathbf{a}$'s reserve price $\mathrm{RP}_{\mathbf{a}} \in \mathbb{R}^+$, temporal discount factor $\delta_{\mathbf{a}} \in (0, 1]$, and deadline $T_{\mathbf{a}} \in \mathbb{N}, T_{\mathbf{a}} > 0$. If the outcome is $FAIL$, $U_{\mathbf{a}}(FAIL) = 0$. The utility function $U_{\mathbf{a}}$ for bargaining outcome (x, t) is defined as:

$$U_{\mathbf{a}}(x,t) = \begin{cases} (\mathrm{RP}_{\mathbf{a}} - x) \cdot \delta^t_{\mathbf{a}} & \text{if } t \leq T_{\mathbf{a}} \text{ and } \mathbf{a} \text{ is a buyer} \\ (x - \mathrm{RP}_{\mathbf{a}}) \cdot \delta^t_{\mathbf{a}} & \text{if } t \leq T_{\mathbf{a}} \text{ and } \mathbf{a} \text{ is a seller} \\ \epsilon < 0 & \text{otherwise} \end{cases}$$

With complete information the appropriate solution concept for the game is the subgame perfect equilibrium in which agents' strategies are in equilibrium in every possible subgame. Such a solution can be found by backward induction. The appropriate solution concept for an extensive-form game with uncertainty is *sequential equilibrium*. A sequential equilibrium is a pair $a = \langle \mu, \sigma \rangle$ (also called

Cite as: Searching for Pure Strategy Equilibria in Bilateral Bargaining With One-sided Uncertainty (Extended Abstract), Bo An, Nicola Gatti and Victor Lesser, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. 1607-1608

an *assessment*) where μ is a belief system that specifies how agents' beliefs evolve during the game and σ specifies agents' strategies. At an equilibrium μ must be *consistent* with respect to σ and σ must be *sequentially rational* given μ.

We assume the one-sided uncertainty regarding the type of the buyer $\mathbf{b}$ (the case of having uncertainty with the type of the seller $\mathbf{s}$ can be analyzed analogously). The buyer $\mathbf{b}$ can be of finitely many types $\{\mathbf{b}_1, \ldots, \mathbf{b}_n\}$ in which buyer $\mathbf{b}_i$ has a reserve price RP_i. The initial belief of $\mathbf{s}$ on $\mathbf{b}$ is $\mu(0) = \langle \Delta^0_{\mathbf{b}}, P^0_{\mathbf{b}} \rangle$ where $\Delta^0_{\mathbf{b}} = \{\mathbf{b}_1, \ldots, \mathbf{b}_n\}$ and $P^0_{\mathbf{b}} = \{\omega^0_{\mathbf{b}_1}, \ldots, \omega^0_{\mathbf{b}_n}\}$ such that $\sum_i \omega^0_{\mathbf{b}_i} = 1$. $\omega^0_{\mathbf{b}_i}$ is the *priori* probability that $\mathbf{b}$ is of type $\mathbf{b}_i$. The belief of $\mathbf{s}$ on the type of $\mathbf{b}$ at time t is $\mu(t)$. The probability assigned by $\mathbf{s}$ to $\mathbf{b} = \mathbf{b}_i$ at time t is denoted $\omega^t_{\mathbf{b}_i}$. Given an assessment $a = \langle \mu, \sigma \rangle$, there are multiple possible bargaining outcomes: outcome $o_{\mathbf{b}_i}$ if $\mathbf{b} = \mathbf{b}_i$. We denote bargaining outcome as $o = \langle o_{\mathbf{b}_1}, \ldots, o_{\mathbf{b}_n} \rangle$.

$\mathbf{s}$'s belief of the type of $\mathbf{b}$ will evolve based on its observed actions and the buyer's equilibrium strategies. On the equilibrium path, $\mathbf{s}$'s belief at any time t is $\mu(t) = \langle \Delta^t_{\mathbf{b}}, P^t_{\mathbf{b}} \rangle$. We need also specify the belief system off the equilibrium path, i.e., when an agent takes an action that is not optimal. We use the *optimistic conjectures* [6]. That is, when $\mathbf{b}$ acts off the equilibrium strategy, agent $\mathbf{s}$ will believe that agent $\mathbf{b}$ is of its "weakest" type, i.e., the type against which the seller would gain the most. This choice is directed to assure the existence of the equilibrium for the largest subset of the space of the parameters. In our case, the weakest type is the buyer type with the highest reserve price.

3. OUR APPROACH

We use the term "choice rule" to characterize buyer types' strategies regarding whether they behave in the same way at a specific decision making point. With pure strategies, buyer types' choice rules are finite. Consider that the belief of $\mathbf{s}$ on the type of $\mathbf{b}$ at time t is $\mu(t) = \Delta_{\mathbf{b}}$ where $|\Delta_{\mathbf{b}}| > 1$ (note that if $|\Delta_{\mathbf{b}}| = 1$, the bargaining from time t becomes the trivial complete information bargaining) and $\iota(t) = \mathbf{b}$. Let the equilibrium offer of buyer type $\mathbf{b}_i \in \Delta_{\mathbf{b}}$ be $x_{\mathbf{b}_i}(t)$. After receiving $\mathbf{b}$'s offer, $\mathbf{s}$ will update its belief and decide whether to accept the offer from $\mathbf{b}$. There are two situations: 1) All buyer types make the same offer. In this case, a *pooling* choice rule is chosen by different buyer types. 2) Buyer types make different offers. That is, a *separating* choice rule is used.

It is easy to see that there are two pooling choice rules depending on whether the seller will accept the offer at time $t+1$: 1) *accepting pooling choice rule* in which all buyer types make the same acceptable offer to $\mathbf{s}$; 2) *rejecting pooling choice rule* in which all buyer types make the same rejectable offer (i.e., -1) to seller $\mathbf{s}$. While the buyer adopts the separating choice rule, some buyer types' equilibrium offers are acceptable to the seller and the number of separating choice rules is drastically reduced due to the following theorem.

We found that there is no equilibrium assessment in pure strategies if buyer types make different acceptable offers at t. Therefore, we only need to consider the following separating choice rules: buyer types $\Delta^a_{\mathbf{b}}$ make an acceptable offer to $\mathbf{s}$ at time t but buyer types $\Delta^r_{\mathbf{b}} = \Delta_{\mathbf{b}} - \Delta^a_{\mathbf{b}}$ make an offer (i.e., -1) that will be rejected by $\mathbf{s}$ at time t. Assume that $\mathbf{b}$ behaves in different ways at an information set $\Delta_{\mathbf{b}}$ at time t where $\Delta_{\mathbf{b}} = \Delta^a_{\mathbf{b}} \cup \Delta^r_{\mathbf{b}}$ at time t. We found that if there is a buyer type $\mathbf{b}_i \in \Delta^a_{\mathbf{b}}$ and a buyer $\mathbf{b}_j \in \Delta^r_{\mathbf{b}}$ such that $\mathrm{RP}_i < \mathrm{RP}_j$, there is no sequential equilibrium for this choice tree. Thus, we only need to consider partitions $\Delta^a_{\mathbf{b}} \cup \Delta^r_{\mathbf{b}} = \Delta_{\mathbf{b}}$ such that for any buyer type $\mathbf{b}_i \in \Delta^a_{\mathbf{b}}$ and any $\mathbf{b}_j \in \Delta^r_{\mathbf{b}}$, $\mathrm{RP}_i > \mathrm{RP}_j$. Thus, the number of separating choice rules is $|\Delta_{\mathbf{b}}| - 1$.

Choice search tree is used to represent agents' choice rules at each decision making point along the bargaining horizon. Each node on a choice search tree is represented as a set of possible types of $\mathbf{b}$ at a time point. Our idea for finding a sequential equilibrium is to search all choice trees in which the buyer's choice rule is clearly specified at each time point. For each choice tree, we use a Bayesian-extension of backward induction to compute agents' sequential equilibrium strategies if it exists. The backward induction starts from all the terminal nodes in the choice tree. The equilibrium strategies of agent $\iota(t)$ at a node at time t depends on agents' equilibrium strategies in the subtree starting from that node. If any condition of equilibrium existence is violated at a node, there is no sequential equilibrium for the choice tree and the backward induction stops.

4. CONCLUSION

While it is very involved to compute sequential equilibria considering all the options at each decision making point, we employ a forward-backward approach: we search forward to find all the choice trees (systems) and we construct backward agents' equilibrium strategies and belief systems for each choice tree. Our approach can be treated as a way of shifting the difficulty of finding a sequential equilibrium in a bargaining game where the buyer has multiple choices to finding a sequential equilibrium in multiple bargaining games in which there is only one choice rule at the buyer's each decision point. To guarantee the completeness of our approach, we enumerate all possible choice trees. By exploiting game theoretic analysis we construct a pair composed of choice rules and belief systems for each possible choice rules. These pairs are parameterized: agents' optimal offers and acceptance at time t depend on the agents' strategies in the following time points till the end of the bargaining. Furthermore, we introduce for each pair some conditions: if they are satisfied, then there is a sequential equilibrium in the subgame starting from time t. For each choice tree, we employ a Bayesian extension of backward induction to derive agents' optimal strategies.

One future research direction is to experimentally evaluate the performance of the derived fully rational equilibrium strategies as compared with heuristics based negotiation strategies. Another future research direction is finding mixed equilibrium strategies for bargaining scenarios in which there is no pure strategy equilibrium. Studying agents' equilibrium strategies for bargaining with two-sided uncertainty is also on the agenda.

5. ACKNOWLEDGMENT

This work was partially supported by the Engineering Research Centers Program of the National Science Foundation under NSF Cooperative Agreement No. EEC-0313747. Any Opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect those of the National Science Foundation.

6. REFERENCES

[1] B. An, N. Gatti, and V. Lesser. Bilateral bargaining with one-sided two-type uncertainty. In *Proc. of the 9th IEEE/WIC/ACM Int. Conf. on Intelligent Agent Technology (IAT)*, pages 403–410, Sep. 2009.

[2] K. Chatterjee and L. Samuelson. Bargaining with two-sided incomplete information: an infinite horizon model with alternating offers. *Review of Economic Studies*, 54(2):175–192, 1987.

[3] N. Gatti, F. D. Giunta, and S. Marino. Alternating-offers bargaining with one-sided uncertain deadlines: an efficient algorithm. *Artificial Intelligence*, 172(8-9):1119–1157, 2008.

[4] P. Miltersen and T. Sorensen. Computing sequential equilibria for two-player games. In *Proc.of the the ACMÍCSIAM Symposium on Discrete Algorithm (SODA)*, pages 107–116, 2006.

[5] A. Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):97–109, 1982.

[6] A. Rubinstein. A bargaining model under incomplete information about time preferences. *Econometrica*, 53(5):1151–1172, 1985.