An Algorithmic Game Theory Framework for Bilateral Bargaining with Uncertainty

(Extended Abstract)

Sofia Ceppi
DEI, Politecnico di Milano, Milano, Italy
ceppi@elet.polimi.it

Nicola Gatti
DEI, Politecnico di Milano, Milano, Italy
ngatti@elet.polimi.it

ABSTRACT

Bilateral bargaining is the most common economic transaction. Customarily, it is formulated as a non-cooperative game with uncertain-information and infinite actions (offers are real-value). Its automation is a long-standing open problem in artificial intelligence and no algorithmic methodology employable regardless of the kind of uncertainty is provided. In this paper, we provide the first step (with one-sided uncertainty) of an algorithmic game theory framework to solve bargaining with any kind of uncertainty. The idea behind our framework is to reduce, by analytical tools, a bargaining problem to a finite game and then to compute, by algorithmic tools, an equilibrium in this game.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Intelligent agents*

General Terms

Algorithms

Keywords

Game theory (cooperative and non-cooperative), Bargaining and negotiation

1. INTRODUCTION

The automation of economic transactions through negotiating software agents is receiving more and more attention in the artificial intelligence community. It is common the idea that autonomous agents can lead to economic contracts more efficient than those drawn up by humans, saving also time and resources. In this paper, we focus on the main bilateral negotiation setting: the *bilateral bargaining*.

A bargaining situation is characterized by the interaction of two agents, a buyer and a seller, who can cooperate to produce a utility surplus by reaching an economic agreement, but they are in conflict on what specific agreement to reach. This is because agents have conflictual interests. A bargaining situation is customarily studied by resorting to game theoretical tools [9], in which each agent is supposed to be rational. The most expressive bilateral bargaining model is the Rubinstein's *alternating-offers* [7]. This protocol prescribes that agents act in turns in alternate fashion and each of them can accept the offer which her opponent made at the previous turn or make a new offer. The utility of the agents over the economic agreements depends on some parameters (i.e., discount factor, deadline, reservation price). In real-world settings, the values of these parameters are private information of the agents. Customarily, agents are assumed to have a probabilistic prior over the values of the opponent.

Solving a bargaining problem means to find the agents' optimal strategies. The alternating-offers is an *infinite-horizon* (agents can indefinitely bargain) *extensive-form* (the game is sequential) *Bayesian* (information is uncertain) game and the number of available actions to each agent is infinite (an offer is a real value). The appropriate solution concept for such a class of games is the *sequential equilibrium* [5]. It is composed of a *belief system*, which describes how agents must update their beliefs during the game, and of *strategies*, which prescribe how rational agents must act.

2. STATE OF THE ART

While solving bargaining with complete-information is easy by using backward induction [9], the study of bargaining with uncertain-information is an open challenging problem. No algorithmic methodology discussed in the literature so far can be applied to this game regardless of the uncertainty kind (i.e., the uncertain parameters) and degree (i.e., the possible values that the parameters can assume). Algorithmic game theory [9] provides general purpose algorithms to search for sequential equilibria [6], but they work only on games with a finite number of actions and they do not produce belief systems off the equilibrium path. This makes such algorithms not suitable for bargaining. Several efforts have been accomplished to extend the backward induction algorithm [2] to solve games with uncertain-information [3]. The basic idea behind these extensions is to break the circularity between strategies and belief system by computing at first the strategies with the initial beliefs and then deriving the beliefs that are consistent with the strategies. However, as shown in [4], the solutions produced by these extensions may not be equilibria, the strategies being not assured to be sequentially rational given the belief system. The microeconomic literature provides some analytical results only for settings without deadlines and with narrow degrees of uncertainty, e.g., over the discount factor of one agents with two possible values [8] and over the reservation price of both agents with two possible values per agent [1]. It

Cite as: An Algorithmic Game Theory Framework for Bilateral Bargaining with Uncertainty (Extended Abstract), S. Ceppi and N. Gatti, *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. 1489-1490

is worth remarking that a large amount of analytical works deal with asymmetric bargaining situations where only one agent makes offer and the other can only accept or reject offers. Finally, it is worth citing an hybrid approach [4] that combines analytical results and searching algorithms to solve the setting in which uncertainty is over the deadline of one agent with an arbitrary number of possible values. This algorithm is proved to be sound and complete and is computationally efficient. However, due to the mathematical machinery it needs to solve a very specific setting of uncertainty, its extension to capture other uncertainty kinds appears to be impractical.

3. THE PROPOSED FRAMEWORK

The aim of the present paper is to provide a framework that can be employed with arbitrary kinds and degrees of uncertainty. Differently from related works, e.g., [3], that focus on searching for equilibria in pure strategies and, in the case there is no pure strategy equilibrium, resort to mixed strategies, we directly search for equilibria in mixed strategies. This is because bargaining with uncertainty may not admit any equilibrium in pure strategies, as shown in [4]. The basic idea behind our algorithm is to solve the bargaining problem by reducing it to a finite game deriving equilibrium strategies such that on the equilibrium path the agents can act only a finite set of actions and then by searching for the agents' optimal strategies on the path. Our framework is structured in the following three steps.

- We analytically derive an assessment $\overline{a} = (\overline{\mu}, \overline{\sigma})$ in which the randomization probabilities of the agents are parameters and such that, when the parameters' values satisfy some conditions, $\overline{a}$ is a strong sequential equilibrium.
- We formulate the problem of finding the values of the agents' randomization probabilities in $\overline{a}$ as the problem of finding a weak sequential equilibrium in a reduced bargaining game with finite actions, and we prove that there always exist values such that $\overline{a}$ is a strong sequential equilibrium.
- We develop an algorithm based on support enumeration to compute an equilibrium in the reduced game and we show that its computational complexity is polynomial in the agents' deadlines.

We apply our framework to settings with one-sided uncertainty (on one agent) over two possible types.

4. THE TWO-TYPE SETTING

In the setting we are studying, buyer's types can have different values of reservation prices (denoted by $RP_{\mathbf{b}}$), temporal discount factor (denoted by $\delta_{\mathbf{b}}$), and temporal deadline (denoted by $T_{\mathbf{b}}$). We call the buyer's types $\mathbf{b}_1$ and $\mathbf{b}_2$. Without loss of generality we assume $T_{\mathbf{b}_1} \leq T_{\mathbf{b}_2}$. We call $\iota(t) : \mathbb{N} \rightarrow \{\mathbf{b}, \mathbf{s}\}$ the player function returning the agent that act at t. We build an assessment $\overline{a}$ such that, on the equilibrium path, the $\iota(t)$'s offers at $t < T_{\mathbf{b}_1}$ belong to a finite set $X(t) \coloneqq \{x^*_{\mathbf{b}_i}(t) : \forall i\}$, where $x^*_{\mathbf{b}_i}(t)$ the $\iota(t)$'s optimal offer at t in the corresponding complete-information game between $\mathbf{b}_i$ and $\mathbf{s}$. Offering at t any $x \notin X(t)$ does not allow $\iota(t)$ to improve her expected utility. In Fig. 1 we show $x^*_{\mathbf{b}_1}(t)$s and $x^*_{\mathbf{b}_2}(t)$s in an example.

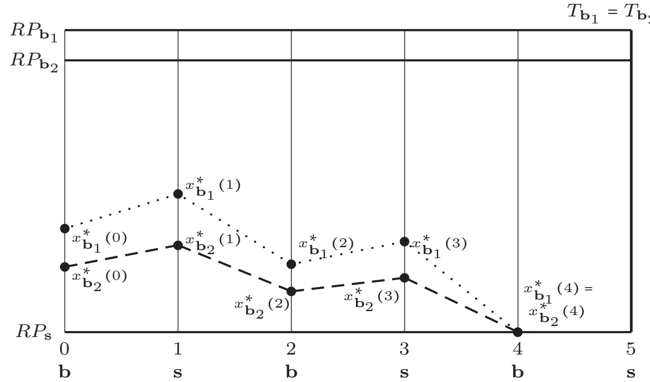

Figure 1: Optimal offers $x^*_{\mathbf{b}_i}(t)$s in the complete information bargaining games between s and $\mathbf{b}_i$.

For each $t < T_{\mathbf{b}_1}$ we rank the values in $X(t)$ in increasing order and we call $\mathbf{b}_s = \arg\min_{i \in \{\mathbf{b}_1, \mathbf{b}_2\}} \{x^*_i(0)\}$ and $\mathbf{b}_w = \arg\max_{i \in \{\mathbf{b}_1, \mathbf{b}_2\}} \{x^*_i(0)\}$ where w means *weak* and s means *strong*. In Fig. 1 we have $\mathbf{b}_w = \mathbf{b}_1$ and $\mathbf{b}_s = \mathbf{b}_2$. The adjectives 'strong' and 'weak' refer to the contractual power of the corresponding buyer's type: in complete-information settings the seller's expected utility is larger when it bargains with $\mathbf{b}_w$ rather than when it bargains with $\mathbf{b}_s$. The basic idea behind $\overline{a}$ is that, when agents are forced to make the offers in $X(t)$, $\mathbf{b}_w$ can gain utility from disguising herself as $\mathbf{b}_s$, making the optimal $\mathbf{b}_s$'s offers, while $\mathbf{b}_s$ prefers to signal her own type, making offers different from the $\mathbf{b}_w$'s ones. That is, $\mathbf{b}_w$ acts in order to increase her expected utility with respect to the situation where $\mathbf{s}$ believes $\mathbf{b}$'s type to be $\mathbf{b}_w$ with certainty.

5. REFERENCES

[1] K. Chatterjee and L. Samuelson. Bargaining under two-sided incomplete information: The unrestricted offers case. *OPER RES*, 36(4):605–618, 1988.

[2] F. Di Giunta and N. Gatti. Bargaining over multiple issues in finite horizon alternating-offers protocol. *ANN MATH ARTIF INTEL*, 47(3-4):251–271, 2006.

[3] S. S. Fatima, M. Wooldridge, and N. R. Jennings. On efficient procedures for multi-issue negotiation. In *TADA-AMEC*, pages 71–84, Hakodate, Japan, May 9 2006.

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

[5] D. R. Kreps and R. Wilson. Sequential equilibria. *ECONOMETRICA*, 50(4):863–894, 1982.

[6] P. B. Miltersen and T. B. Sorensen. Computing sequential equilibria for two-player games. In *ACM-SIAM SODA*, pages 107–116, 2006.

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

[8] A. Rubinstein. A bargaining model with incomplete information about time preferences. *ECONOMETRICA*, 53(5):1151–1172, 1985.

[9] Y. Shoham and K. Leyton-Brown. *Multiagent Systems: Algorithmic, Game Theoretic and Logical Foundations.* 2008.