Extending Alternating-Offers Bargaining in One-to-Many and Many-to-Many Settings

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

Automating negotiations in markets where multiple buyers and sellers operate is a scientific challenge of extraordinary importance. One-to-one negotiations are classically studied as bilateral bargaining problems, while one-to-many and many-to-many negotiations are studied as auctioning problems. This paper aims at bridging together these two approaches, analyzing agents' strategic behavior in one-to-many and many-to-many negotiations when agents follow the alternating-offers bargaining protocol [5]. First, we propose a novel mechanism that captures the peculiarities of these settings. Then, we preliminarily explore how uncertainty over reserve prices and deadlines can affect equilibrium strategies. Surprisingly, the computation of the equilibrium for realistic ranges of the parameters in one-to-many settings is reduced to the computation of the equilibrium either in one-to-one settings with uncertainty or in one-to-many settings without uncertainty.

1. Introduction

The focus of this work is on analyzing agents' strategic behavior in one-to-many and many-to-many negotiations in which agents are negotiating with multiple trading partners and, at the same time, are facing competition from competitors. The subgame perfect equilibrium is presented and equilibrium properties, such us uniqueness, are discussed. We also provide a preliminary extension to the incomplete information setting. The main goal of this paper is to begin to understand which factors are affecting agents' bargaining position relative to others and agents' equilibrium bargaining strategies. This analysis is designed to provide some suggestions for designing negotiation agents in an electronic marketplace in which agents often are involved in many-to-many negotiations. While there has been much experimental work (e.g., [2], [4]) on one-to-many and many-to-many negotiations, to our best knowledge, this paper is the first work to provide a game theoretical analysis of agents' strategic interactions in concurrent negotiations.

This paper analyzes agents' strategic behavior in concurrent one-to-many negotiation and many-to-many negotiation. For complete information settings, the computational complexity when there are many buyers and many sellers in our protocol is essentially the same in the situation where the negotiations are purely bilateral. We preliminarily explore how uncertainty over reserve prices and deadlines can affect equilibrium strategies. We observe that agents' bargaining power are affected by the proposing ordering and market competition. We also find that the computation of the equilibrium for realistic ranges of the parameters in one-to-many settings reduces to the computation of the equilibrium either in one-to-one settings with uncertainty or in one-to-many settings without uncertainty.

The rest of this paper proceeds as follows: We start with bilateral negotiation in Section 2. Section 3 discusses one-to-many negotiation and Section 4 investigates many-to-many negotiation. Section 5 concludes this paper. The proofs of theorems and lemmas are reported in [1].

2. Bilateral Negotiation

We study a discrete time bilateral negotiation between a buyer $\mathbf{b}$ and a seller $\mathbf{s}$. A finite horizon 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 $\iota : \mathbb{N} \to \{\mathbf{b}, \mathbf{s}\}$ 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 time $t > 0$ are: 1) $\mathit{offer}[x]$, where $x \in \mathbb{R}$ is the proposed price for the good; 2) exit, which implies that negotiation between $\mathbf{b}$ and $\mathbf{s}$ fails; and 3) accept, which implies that $\mathbf{b}$ and $\mathbf{s}$ make 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 $FAIL$.

Each agent $\mathbf{a} \in \{\mathbf{b}, \mathbf{s}\}$ has a utility function $U_\mathbf{a} : (\mathbb{R} \times \mathbb{N}) \cup FAIL \to \mathbb{R}$, which represents its gain over the bargaining outcomes. Each utility function $U_\mathbf{a}$ depends on $\mathbf{a}$'s reserve price $RP_\mathbf{a} \in \mathbb{R}^+$, temporal discount factor $\delta_\mathbf{a} \in (0, 1]$, and deadline $T_\mathbf{a} \in \mathbb{N}, T_\mathbf{a} > 0$. For ease of analysis, we assume that agents have different reserve prices throughout this paper.

The utility function $U_\mathbf{a}$ for bargaining outcome (x, t) is:

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

If the outcome is $FAIL$, then $U_\mathbf{a}(FAIL) = 0$. Notice that the assignment of a strictly negative value ($\epsilon > 0$) to $U_\mathbf{a}$ after agent $\mathbf{a}$'s deadline implies that an agent, after its deadline, strictly prefers to exit the negotiation.

Initially, it is determined the time point T where the game rationally stops: it is $T = \min(T_\mathbf{b}, T_\mathbf{s})$. The equilibrium outcome of every subgame starting from $t \geq T$ is $FAIL$, since at least one agent will make $exit$. Therefore, at $t = T$ agent $\iota(T)$ would accept any offer x which gives it a utility not worse than $FAIL$, namely, any offer x such that $U_{\iota(T)}(x, T) \geq 0$. From $t = T-1$ back to $t = 0$ it is possible to find the optimal offer agent $\iota(t)$ can make at t, if it makes an offer, and the offers that it would accept. $x^*(t)$ denotes the optimal offer of agent $\iota(t)$ at t. $x^*(t)$ is the offer such that, if $t < T-1$, agent $\iota(t+1)$ is indifferent at $t+1$ between accepting it and rejecting it to make its optimal offer $x^*(t+1)$ and, if $t = T-1$, agent $\iota(t+1)$ is indifferent at $t+1$ between accepting it and exiting from negotiation. Formally, $x^*(t)$ is such that $U_{\iota(t+1)}(x^*(t)), t) = U_{\iota(t+1)}(x^*(t+1)$, t+1) if $t < T-1$ and $U_{\iota(t+1)}(x^*(t), t) = 0$ if $t = T-1$. The offers agent $\iota(t)$ would accept at t are all those offers that give it a utility no worse than the utility given by offering $x^*(t)$. The equilibrium strategy of any subgame starting from $t < T$ prescribes that agent $\iota(t)$ offers $x^*(t)$ at t and agent $\iota(t+1)$ accepts it.

Backward propagation is used to provide a recursive formula for $x^*(t)$: given value x and agent $\mathbf{a}$, we call backward propagation of value x for agent $\mathbf{a}$ the value y such that $U_{\mathbf{a}}(y, t-1) = U_{\mathbf{a}}(x, t)$; we employ the arrow notation $x_{\leftarrow \mathbf{a}}$ for backward propagations. Formally, $x_{\leftarrow \mathbf{b}} = RP_{\mathbf{b}} - (RP_{\mathbf{b}} - x) \cdot \delta_{\mathbf{b}}$ and $x_{\leftarrow \mathbf{s}} = RP_{\mathbf{s}} + (x - RP_{\mathbf{s}}) \cdot \delta_{\mathbf{s}}$. If a value x is backward propagated n times for agent $\mathbf{a}$, we write $x_{\leftarrow n[\mathbf{a}]}$, e.g. $x_{\leftarrow 2[\mathbf{a}]} = (x_{\leftarrow \mathbf{a}})_{\leftarrow \mathbf{a}}$. If a value is backward propagated for more than one agent, we list them left to right in the subscript, e.g., $x_{\leftarrow \mathbf{b} 2[\mathbf{s}]} = ((x_{\leftarrow \mathbf{b}})_{\leftarrow \mathbf{s}})_{\leftarrow \mathbf{s}}$. The values of $x^*(t)$ can be calculated recursively from $t = T-1$ back to $t = 0$ as

$$x^*(t) = \begin{cases} RP_{\iota(t+1)} & \text{if } t = T-1 \\ (x^*(t+1))_{\leftarrow \iota(t+1)} & \text{if } t < T-1 \end{cases}$$

It can be easily observed that $x_{\leftarrow \mathbf{b}} \geq x$ as $x_{\leftarrow \mathbf{b}} - x = RP_{\mathbf{b}} - (RP_{\mathbf{b}} - x) \cdot \delta_{\mathbf{b}} - x = (1 - \delta_{\mathbf{b}})(RP_{\mathbf{b}} - x) \geq 0$, and $x_{\leftarrow \mathbf{s}} \leq x$ as $x_{\leftarrow \mathbf{s}} - x = RP_{\mathbf{s}} + (x - RP_{\mathbf{s}}) \cdot \delta_{\mathbf{s}} - x = (\delta_{\mathbf{s}} - 1)(x - RP_{\mathbf{s}}) \leq 0$. In addition, if $x \leq RP_{\mathbf{b}}$, it follows that $x_{\leftarrow \mathbf{b}} \leq RP_{\mathbf{b}}$. Similarly, If $x \geq RP_{\mathbf{s}}$, $x_{\leftarrow \mathbf{s}} \geq RP_{\mathbf{s}}$.

Finally, the equilibrium strategies of $\mathbf{b}$ can be defined as (the equilibrium strategies of $\mathbf{s}$ can be defined analogously)

$$\sigma^*_{\mathbf{b}}(t) = \begin{cases} t = 0 & \textit{offer}[x^*(0)] \\ 0 < t < T & \begin{cases} \text{if } \sigma_{\mathbf{s}}(t-1) = \textit{offer}[x] \text{ with } x \leq x^*(t)_{\leftarrow \mathbf{b}} & \textit{accept} \\ \text{otherwise} & \textit{offer}[x^*(t)] \end{cases} \\ T \leq t \leq T_{\mathbf{b}} & \begin{cases} \text{if } \sigma_{\mathbf{s}}(t-1) = \textit{offer}[x] \text{ with } x \leq RP_{\mathbf{b}} & \textit{accept} \\ \text{otherwise} & \textit{exit} \end{cases} \\ T_{\mathbf{b}} < t & \textit{exit} \end{cases}$$

Therefore, at equilibrium, the two agents will reach an agreement at time $t = 1$ and the agreement price is $x^*(0)$.

3. One-to-many Alternating-Offers Negotiation

3.1. Negotiation Mechanism

In this section, we consider the situation where there is one buyer agent $\mathbf{b}$ and a set $\mathcal{S} = \{\mathbf{s}_1, \ldots, \mathbf{s}_n\}$ of n seller agents. As in [2], [4], a buyer synchronously negotiates with multiple sellers in discrete time. We use the term "negotiation thread" for the bargaining between $\mathbf{b}$ and a seller $\mathbf{s}_i$ and we denote it by $\Im_{\mathbf{b},\mathbf{s}_i}$. Furthermore, we denote by $\iota(\Im_{\mathbf{b},\mathbf{s}_i}, t)$ the agent that acts at t in the negotiation thread $\Im_{\mathbf{b},\mathbf{s}_i}$. We assume that if $\iota(\Im_{\mathbf{b},\mathbf{s}_i}, t) = \mathbf{b}$ then $\iota(\Im_{\mathbf{b},\mathbf{s}_j}, t) = \mathbf{b}$ for all j. That is, $\mathbf{b}$ simultaneously acts in all threads. Therefore, if $\mathbf{b}$ is proposing at time t, $\iota(t) = \mathbf{b}$. Otherwise, $\iota(t) = \mathcal{S}$.

We modify the alternating-offers mechanism by introducing an action *confirm* to avoid agents' non-reasonable behaviors. The sellers' action space is $A = \{\textit{offer}[x], \textit{accept}, \textit{exit}, \textit{confirm}\}$, whereas the buyer's action space is the Cartesian product $\times_{i=1}^{n} A$. Legal actions for the buyer are all the pure strategies $\sigma_{\mathbf{b}} = \langle \sigma_{\mathbf{b},\mathbf{s}_1}, \ldots, \sigma_{\mathbf{b},\mathbf{s}_n} \rangle$ such that: if $\sigma_{\mathbf{s}_i}(t-1) \neq \textit{accept}$, then $\sigma_{\mathbf{b},\mathbf{s}_i}(t) \in \{\textit{offer}[x], \textit{accept}, \textit{exit}\}$ except when $t = 0$, *accept* is not available, otherwise $\sigma_{\mathbf{b},\mathbf{s}_i}(t) \in \{\textit{confirm}, \textit{exit}\}$. Legal actions for the sellers are defined analogously: if $\sigma_{\mathbf{b},\mathbf{s}_i}(t-1) \neq \textit{accept}$, then $\sigma_{\mathbf{s}_i}(t) \in \{\textit{offer}[x], \textit{accept}, \textit{exit}\}$ except when $t = 0$, *accept* is not available, otherwise $\sigma_{\mathbf{s}_i}(t) \in \{\textit{confirm}, \textit{exit}\}$. The action *confirm* is allowed only after making the action $accept$.

The outcome of a single negotiation thread $\Im_{\mathbf{b},\mathbf{s}_i}$ is $FAIL$ if either $\mathbf{b}$ or $\mathbf{s}_i$ exits, whereas it is an agreement (x, t) if $\sigma_{\iota(\Im_{\mathbf{b},\mathbf{s}_i}, t)}(t) = \textit{confirm}$, where x is such that $\sigma_{\iota(\Im_{\mathbf{b},\mathbf{s}_i}, t-2)}(t-2) = \textit{offer}[x]$. Notice that, in absence of the action *confirm*, if $\mathbf{b}$ makes offers to more sellers and all these accept, $\mathbf{b}$ must buy more items. In presence of the action *confirm*, $\mathbf{b}$ is in the position to choose only one contract. Thus, the following process is needed for implementing an agreement: one agent proposes a price, the other agent accepts the offer, then the first agent confirms the contract.

We redefine $\mathbf{b}$'s utility as follows. If $\mathbf{b}$ has reached more than one agreement, let (x_{first}, t_{first}) be the agreement such that, for any other agreement (x_j, t_j), (1) $t_{first} \leq t_j$ and (2) $x_{first} \leq x_j$ if $t_{first} = t_j$. Let i_{first} be the seller involved in the agreement (x_{first}, t_{first}). Agent $\mathbf{b}$'s utility is defined over the set of agreements it reached:

$$U_{\mathbf{b}}(\{(x_i, t_i)\}) = \begin{cases} (RP_{\mathbf{b}} - x_{first}) \cdot \delta_{\mathbf{b}}^{t_{first}} - \sum_{j \neq i_{first}} x_j & \text{if } t_{first} \leq T_{\mathbf{b}} \\ -\epsilon & \text{otherwise} \end{cases}$$

That is, $\mathbf{b}$ receives a positive utility from the first agreement, whereas all the other agreements reduce $\mathbf{b}$'s utility.

3.2. Agents' Equilibrium Strategies

Let $\mathcal{S}_{=t}$ be the set of sellers whose deadline is t, i.e., $\mathcal{S}_{=t} = \{\mathbf{s}_i | T_{\mathbf{s}_i} = t\}$. Let $\mathcal{S}_t$ be the set of sellers which have no shorter deadline than t, i.e., $\mathcal{S}_t = \{\mathbf{s}_i | T_{\mathbf{s}_i} \geq t\} = \cup_{t' \geq t} \mathcal{S}_{=t}$. We assume that the sellers $\mathcal{S}_t$ are ranked according to their reserve prices. We denote by $\mathcal{S}^i_t$ ($\mathcal{S}^i_{=t}$) the seller with the i^{th} lowest reserve price in $\mathcal{S}_t$ ($\mathcal{S}_{=t}$). Let $x^*_{\mathbf{b},\mathbf{s}_i}(t)$ be $\mathbf{b}$'s optimal offer to $\mathbf{s}_i$ at time t if $\iota(\Im_{\mathbf{b},\mathbf{s}_i}, t) = \mathbf{b}$ and $x^*_{\mathbf{s}_i,\mathbf{b}}(t)$ be $\mathbf{s}_i$'s optimal offer to agent $\mathbf{b}$ at time t if $\iota(\Im_{\mathbf{b},\mathbf{s}_i}, t) = \mathbf{s}_i$.

The negotiation deadline for the negotiation thread between $\mathbf{b}$ and $\mathbf{s}_i$ is $T_{\mathbf{b},\mathbf{s}_i} = \min(T_{\mathbf{b}}, T_{\mathbf{s}_i})$. After $T_{\mathbf{b},\mathbf{s}_i}$, at least one agent will have no interest in reaching agreements. Obviously, the negotiation deadline for $\mathbf{b}$ is $T = \max_{\mathbf{s}_i \in \mathcal{S}}\{T_{\mathbf{b},\mathbf{s}_i}\}$.

Lemma 1: It is $\mathbf{b}$'s weakly dominant strategy to make the same offer to all the sellers in $\mathcal{S}_{t+2}$ at each time t.

According to Lemma 1 we can assume $x^*_{\mathbf{b},\mathbf{s}_i}(t) = x^*_{\mathbf{b},\mathbf{s}_j}(t)$ for all i, j. For simplicity, we denote such offer by $x^*_{\mathbf{b}}(t)$.

Theorem 2: In the one-to-many negotiation, the sequences of equilibrium offers $x^*_{\mathbf{b}}$ and $x^*_{\mathbf{s}_i}$ are:

$$x^*_{\mathbf{b}}(t) = \begin{cases} RP_{\mathcal{S}^1_{t+2}} & t = T-2 \text{ or } t = T_{\mathcal{S}^1_{t+2}} - 2 \\ \min\{x^*_{\mathcal{S}^1_{t+2}}(t+1)_{\leftarrow \mathcal{S}^1_{t+2}}, RP_{\mathcal{S}^2_{t+2}}\} & t < T-2 \text{ and } t \neq T_{\mathcal{S}^1_{t+2}} - 2 \end{cases},$$

$$x^*_{\mathbf{s}_i}(t) = \begin{cases} \max\{RP_{\mathbf{s}_i}, RP_{\mathcal{S}^2_T}\} & t = T-2 \\ \max\{RP_{\mathbf{s}_i}, \min\{RP_{\mathcal{S}^2_{t+2}}, x^*_{\mathbf{b}}(t+1)_{\leftarrow \mathbf{b}}\}\} & t < T-2 \end{cases}.$$

Agent's equilibrium strategies are similar to those discussed in Section 2, but $\sigma_{\mathbf{b},\mathbf{s}_i}$ prescribes that:

- $\mathbf{b}$ accepts the offer x made by $\mathbf{s}_i$ at t if: $x \leq x^*_{\mathbf{b}}(t)_{\leftarrow \mathbf{b}}$ and x is the lowest received offer. If more than one seller has offered x, than $\mathbf{b}$ accepts the offer made by the seller with the lowest reserve price;
- $\mathbf{b}$ confirms an accept of $\mathbf{s}_i$ at t if: $\sigma_{\mathbf{b}}(t-2) = \mathit{offer}[x]$ with $x \leq x^*_{\mathbf{b}}(t)_{\leftarrow 2[\mathbf{b}]}$ and, among all the sellers that have accepted $\sigma_{\mathbf{b}}(t-2)$, $\mathbf{s}_i$ is the one with the lowest reserve price;

and $\sigma_{\mathbf{b},\mathbf{s}_i}$ prescribes that:

- $\mathbf{s}_i$ confirms the accept of $\mathbf{b}$ at t if: $\sigma_{\mathbf{s}_i}(t-2) = \mathit{offer}[x]$ with $x \geq \max\{x^*_{\mathbf{s}_i}(t)_{\leftarrow 2[\mathbf{s}_i]}, RP_{\mathbf{s}_i}\}$.

The computational complexity of the backward induction is $\mathcal{O}(nT)$ as it will go through all the time points and at each time point, each agent has at most three possible optimal actions. The equilibrium agreement is reached at $t = 2$ between $\mathbf{b}$ and $\mathcal{S}^1_2$ and it is $(x^*_{\mathbf{b}}(0), 2)$ if $\iota(0) = \mathbf{b}$ and $(x^*_{\mathcal{S}^1_2}(0), 2)$ otherwise. It can be easily observed that $RP_{\mathcal{S}^1_2} \leq x^*_{\mathbf{b}}(0), x^*_{\mathcal{S}^1_2}(0) \leq RP_{\mathcal{S}^2_2}$. The result about agreement price is intuitive in the following sense: obviously, the agreement price cannot be lower than each seller's reserve price. But it also cannot be higher than the second lowest price as, if so, there is at least another seller who is willing to sell for less and make an agreement with the buyer. Let us remark an observation. Consider the situation wherein $\iota(0) = \mathcal{S}$ and $x^*_{\mathcal{S}^1_2} = RP_{\mathcal{S}^2_2}$. Although both $\mathcal{S}^1_2$ and $\mathcal{S}^2_2$ have the same equilibrium offer, i.e., $RP_{\mathcal{S}^2_2}$, the equilibrium strategy of $\mathbf{b}$ prescribes that $\mathbf{b}$ must accept only the offer made by $\mathcal{S}^1_2$. In the case $\mathbf{b}$ accepts the offer by $\mathcal{S}^2_2$ or randomizes over accepting those offers, $\mathcal{S}^1_2$'s optimal action at $t = 0$ does not exist, being $\lim_{\varepsilon \to 0}(\mathcal{S}^2_2 - \varepsilon)$ with $\varepsilon \neq 0$.

Theorem 3: Agents' strategies on the equilibrium path are unique except when $RP_{\mathcal{S}^1_2} = RP_{\mathbf{s}_i}$ for more than one i.

When the reserve price of more sellers is equal to $RP_{\mathcal{S}^1_2}$, all these sellers will offer their reserve price and $\mathbf{b}$ can accept any single offer among these. However, all the equilibria are equivalent in terms of agents' payoffs. As we assume that agents have different reserve prices, the equilibrium is unique.

3.3. Incomplete information

Although agents' equilibrium strategies depend on the values of the parameters of all the agents, for a large subset of the space of the parameters the equilibrium outcome depends on the values of a narrow number of parameters.

Theorem 4: When 1) $T_{\mathcal{S}^2_2} > 2$ if $\iota(0) = \mathbf{b}$ and 2) $(RP_{\mathbf{s}})_{\leftarrow \mathcal{S}^1_2 \mathbf{b}} \geq RP_{\mathbf{s}}$ for any seller $\mathbf{s} \in \mathcal{S}$, the equilibrium outcome depends only on the parameters of $\mathbf{b}$ (i.e., $RP_{\mathbf{b}}$, $\delta_{\mathbf{b}}$, $T_{\mathbf{b}}$), $\mathcal{S}^1_2$ (i.e., $RP_{\mathcal{S}^1_2}$, $\delta_{\mathcal{S}^1_2}$, $T_{\mathcal{S}^1_2}$), and on the reserve price $RP_{\mathcal{S}^2_2}$ of $\mathcal{S}^2_2$. In these situations the equilibrium outcome can be produced as follows:

1) find the sequence of the optimal offers under the assumption that $\mathcal{S}^1_2$ is the unique seller, say $y(t)$, and
2) assign $x^*_{\mathbf{b}}(0) = \min\{y(0), (RP_{\mathcal{S}^2_2})_{\leftarrow \mathcal{S}^1_2}\}$ if $\iota(0) = \mathbf{b}$ and assign $x^*_{\mathcal{S}^1_2}(0) = \min\{y(0), RP_{\mathcal{S}^2_2}\}$ if $\iota(0) = \mathcal{S}$.

This is to say that the equilibrium outcome does not depend on the values of $\delta_{\mathcal{S}^2_2}$, $T_{\mathcal{S}^2_2}$, and on the parameters of all the other sellers. This is of paramount importance since complex settings with a high degree of uncertainty can be easily solved when 1) $T_{\mathcal{S}^2_2} > 2$ if $\iota(0) = \mathbf{b}$ and 2) $(RP_{\mathbf{s}})_{\leftarrow \mathcal{S}^1_2 \mathbf{b}} \geq RP_{\mathbf{s}}$ for any seller $\mathbf{s} \in \mathcal{S}$. Indeed, the above algorithm produces the equilibrium outcome even when $\delta_{\mathcal{S}^i_2}$ with $i > 1$, $T_{\mathcal{S}^i_2}$ with $i > 1$, and $RP_{\mathcal{S}^i_2}$ with $i > 2$ are uncertain. We can write condition $(RP_{\mathcal{S}^2_2})_{\leftarrow \mathcal{S}^1_2 \mathbf{b}} \geq RP_{\mathcal{S}^2_2}$ as

$$(RP_{\mathbf{b}} - RP_{\mathcal{S}^1_2}) \geq (RP_{\mathcal{S}^2_2} - RP_{\mathcal{S}^1_2}) \frac{1 - \delta_{\mathbf{b}} \delta_{\mathcal{S}^1_2}}{1 - \delta_{\mathbf{b}}}.$$

It can be easily observed that, in common settings where $RP_{\mathbf{b}} \gg RP_{\mathcal{S}^2_2}$ and $\delta_{\mathcal{S}^1_2}$ is close to 1, the above condition is satisfied.

Now, we focus on the uncertainty over $\mathbf{b}$'s and $\mathcal{S}^1_2$'s parameters. The values of these parameters affect the equilibrium outcome and therefore in presence of uncertainty over them we need to compute agents' equilibrium strategies to derive the equilibrium outcome. Currently, the literature provides algorithms to compute agents' equilibrium strategies only in bilateral settings without outside option with one-sided uncertainty over deadlines [3]. Since the number of available actions is infinite, no algorithms such as Lemke-Howson [6] can be employed to compute a sequential equilibrium.

When $RP_{\mathcal{S}^2_2} \leq (RP_{\mathcal{S}^2_2})_{\leftarrow \mathcal{S}^1_2 \mathbf{b}}$, the algorithm presented in [3] can be easily extended to capture uncertainty in one-to-many bargaining. More precisely, we have that:

- when $T_{\mathbf{b}}$ is uncertain, whereas $T_{\mathcal{S}^1_2}$ is certain, then agents' equilibrium strategies can be produced by employing the algorithm presented in [3] where the buyer is $\mathbf{b}$ and the seller is $\mathcal{S}^1_2$ and upper bounding the optimal offers to $RP_{\mathcal{S}^2_2}$ if $\iota(0) = \mathbf{b}$ and to $(RP_{\mathcal{S}^2_2})_{\leftarrow \mathcal{S}^1_2}$ if $\iota(t) = \mathcal{S}$;
- when $T_{\mathcal{S}^1_2}$ is uncertain, whereas $T_{\mathbf{b}}$ is certain, then agents' equilibrium strategies can be computed.

Settings with a higher degree of uncertainty, such as when both $T_{\mathbf{b}}$ and $T_{\mathcal{S}^1_2}$ are uncertain, need further exploration.

4. Many-to-Many Setting Analysis

4.1. Negotiation Mechanism

In this section, we propose a bargaining model for many-to-many negotiation where m buyer agents $\mathcal{B} = \{\mathbf{b}_1, \ldots, \mathbf{b}_m\}$ negotiate n seller agents $\mathcal{S} = \{\mathbf{s}_1, \ldots, \mathbf{s}_n\}$. In this case, each agent concurrently negotiates with many trading partners. Agent $\mathbf{b}_j$'s concurrent negotiation includes at most n threads $\Im_{\mathbf{b}_j,\mathcal{S}} = \{\Im_{\mathbf{b}_j,\mathbf{s}_i} | \mathbf{s}_i \in \mathcal{S}\}$, where $\Im_{\mathbf{b}_j,\mathbf{s}_i}$ represents the negotiation thread between $\mathbf{b}_j$ and seller $\mathbf{s}_i$. We still assume that, at each time, either the buyers propose to all the sellers ($\iota(t) = \mathcal{B}$) or the sellers propose to all the buyers ($\iota(t) = \mathcal{S}$). Similarly, let $\mathcal{B}_{=t}$ be the set of buyers whose deadline is t, i.e., $\mathcal{B}_{=t} = \{\mathbf{b}_j | T_{\mathbf{b}_j} = t\}$. Let $\mathcal{B}_t$ be the set of buyers which have no shorter deadline than t and $\mathcal{B}_t^i$ ($\mathcal{B}_{=t}^i$) is the buyer with the i^{th} highest reserve price in $\mathcal{B}_t$ ($\mathcal{B}_{=t}$).

We still use action *confirm* to avoid one agent's making more than one final agreement. Buyers and sellers' action space and agents' legal actions at each time are the same as that in one-to-many negotiation. The utility functions of the buyer agents are exactly those defined in the previous section. However, we need to refine the utility function of $\mathbf{s}_i$ as it can potentially sell more items, but it has only one item to sell. We redefine $\mathbf{s}_i$'s utility as follows. If $\mathbf{s}_i$ has reached more than one final agreement, it gets a utility of $-\infty$.

4.2. Agents' Equilibrium Strategies

The negotiation deadline for the negotiation between $\mathbf{b}_j$ and $\mathbf{s}_i$ is $T_{\mathbf{b}_j,\mathbf{s}_i} = \min(T_{\mathbf{b}_j}, T_{\mathbf{s}_i})$. The negotiation deadline for the agent $\mathbf{b}_j$ is $T_{\mathbf{b}_j,\mathcal{S}} = \max_{\mathbf{s}_i \in \mathcal{S}} T_{\mathbf{b}_j,\mathbf{s}_i}$. Let $x^*_{\mathbf{b}_j,\mathbf{s}_i}(t)$ be $\mathbf{b}_j$'s optimal offer to agent $\mathbf{s}_i$ at t if $\iota(t) = \mathcal{B}$ and $x^*_{\mathbf{s}_i,\mathbf{b}_j}(t)$ be $\mathbf{s}_i$'s optimal offer to agent $\mathbf{b}_j$ at time t if $\iota(t) = \mathcal{S}$.

Lemma 5: It is each agent's dominant strategy to propose the same price to all the trading partners at each time t.

Then we use $x^*_{\mathbf{b}_j}(t)$ for short to represent $\mathbf{b}_j$'s optimal offer at t if $\iota(t) = \mathcal{B}$ and use $x^*_{\mathbf{s}_i}(t)$ to represent $\mathbf{s}_i$'s optimal offer at time t if $\iota(t) = \mathcal{S}$.

Lemma 6: In equilibrium, agents of the same type should have the same equilibrium winning price (a price acceptable to agents of the different type).

Theorem 7: In the many-to-many negotiation, the sequences of optimal offers in equilibrium are: Buyer $\mathbf{b}_j$'s optimal offer at time $t \leq T_{\mathbf{b}_j} - 2$ is $x^*_{\mathbf{b}_j}(t) = \min(RP_{\mathbf{b}_j}, x^*_{\mathcal{B}}(t))$. Seller $\mathbf{s}_i$'s optimal offer at $t \leq T_{\mathbf{s}_i} - 2$ is $x^*_{\mathbf{s}_i}(t) = \max(RP_{\mathbf{s}_i}, x^*_{\mathcal{S}}(t))$.

$x^*_{\mathcal{B}}(t)$ is given by: 1) At $t = T-2$, $x^*_{\mathcal{B}}(t) = RP_{\mathcal{S}_{t+2}^{|\mathcal{B}_{t+2}|}}$ if $|\mathcal{B}_{t+2}| \leq |\mathcal{S}_{t+2}|$; otherwise, $x^*_{\mathcal{B}}(t) = RP_{\mathcal{B}_{t+2}^{|\mathcal{S}_{t+2}|+1}}$. 2) At $t < T-2$, $x^*_{\mathcal{B}}(t) = \max\left\{RP_{\mathcal{B}_{t+2}^{|\mathcal{S}_{t+2}|+1}}, \left\{\{x^*_{\mathbf{s}_i}(t+1)_{\leftarrow \mathbf{s}_i} | \mathbf{s}_i \in \mathcal{S}_{t+3}\} \cup \{RP_{\mathbf{s}_i} | \mathbf{s}_i \in \mathcal{S}_{t+2} - \mathcal{S}_{t+3}\}\right\}_{|\mathcal{S}_{t+2}|}\right\}$ if $|\mathcal{S}_{t+2}| < |\mathcal{B}_{t+2}|$. Otherwise, $x^*_{\mathcal{B}}(t) = \left\{\{x^*_{\mathbf{s}_i}(t+1)_{\leftarrow \mathbf{s}_i} | \mathbf{s}_i \in \mathcal{S}_{t+3}\} \cup \{RP_{\mathbf{s}_i} | \mathbf{s}_i \in \mathcal{S}_{t+2} - \mathcal{S}_{t+3}\}\right\}_{|\mathcal{B}_{t+2}|}$. In the above equations, $\mathcal{Y}_i$ ($\mathcal{Y}^i$) denotes the i^{th} smallest (largest) value in the value set $\mathcal{Y}$.

$x^*_{\mathcal{S}}(t)$ is given by: 1) At $t = T-2$, $x^*_{\mathcal{S}}(t) = RP_{\mathcal{B}_T^{|\mathcal{S}_{t+2}|}}$ if $|\mathcal{S}_T| \leq |\mathcal{B}_T|$, $x^*_{\mathcal{S}}(t) = RP_{\mathcal{S}_{t+2}^{|\mathcal{B}_{t+2}|+1}}$ if $|\mathcal{S}_{t+2}| > |\mathcal{B}_{t+2}|$. 2) At $t < T-2$, $x^*_{\mathcal{S}}(t) = \left\{\{x^*_{\mathbf{b}_j}(t+1)_{\leftarrow \mathbf{b}_j} | \mathbf{b}_j \in \mathcal{B}_{t+3}\} \cup \{RP_{\mathbf{b}_j} | \mathbf{b}_j \in \mathcal{B}_{t+2} - \mathcal{B}_{t+3}\}\right\}^{|\mathcal{S}_{t+2}|}$ if $|\mathcal{S}_{t+2}| \leq |\mathcal{B}_{t+2}|$. Otherwise, $x^*_{\mathcal{S}}(t) = \min\left\{RP_{\mathcal{S}_{t+2}^{|\mathcal{B}_{t+2}|+1}}, \left\{\{x^*_{\mathbf{b}_j}(t+1)_{\leftarrow \mathbf{b}_j} | \mathbf{b}_j \in \mathcal{B}_{t+3}\} \cup \{RP_{\mathbf{b}_j} | \mathbf{b}_j \in \mathcal{B}_{t+2} - \mathcal{B}_{t+3}\}\right\}^{|\mathcal{B}_{t+2}|}\right\}$.

Based on $x^*_{\mathbf{b}_j}(t)$ and $x^*_{\mathbf{s}_i}(t)$, we can get agents' optimal actions in the same way as that in Theorem 2 except that an agent needs to use the following rule while accepting offers or confirming accepts: a buyer $\mathbf{b}_j$ accepts the offer x made by $\mathbf{s}_i$ at t if: $x \leq x^*_{\mathbf{b}_j}(t)_{\leftarrow \mathbf{b}_j}$ and x is the lowest received offer. If more than one seller has offered x and buyer $\mathbf{b}_j$ has the q^{th} highest reserve price in $\mathcal{B}_{t+2}$, $\mathbf{b}_j$ accepts the offer made by the seller with the q^{th} lowest reserve price in sellers $\mathcal{S}_{t+2}$.[1] Similarly, if buyer $\mathbf{b}_j$ intends to confirm an agreement with price x and multiple sellers have made the same agreement, $\mathbf{b}_j$ will confirm the agreement made by the seller with the q^{th} lowest reserve price in sellers $\mathcal{S}_{t+2}$. To save space, the details of sellers' optimal actions are omitted.

5. Conclusion

One major motivation of the study of bargaining theory is designing successful bargaining agents in practical dynamic markets. Although constraints, complexity, and uncertainty make it impractical to develop optimal negotiation strategies, this paper provides some useful guidelines for designing negotiation agents. For example, market competition plays a central role in deciding the market equilibrium, agents need to make the same offer to all the trading partners at each time.

One future research direction is theoretically analyzing agents' strategic behavior in many-to-many negotiation when agents have incomplete information about agents' reserve prices, and discounting factors. It would be interesting to investigate an agent's incentive to misrepresent its preference in a market where a single agent's influence on the market equilibrium will decrease with the increase of the scale of the market.

6. Acknowledgement

This work was supported in part 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.

References

[1] B. An, N. Gatti, and V. Lesser. Extending alternating-offers bargaining in one-to-many and many-to-many settings. Technical report, University of Massachusetts, Amherst (UMASS-CS-2009-016), 2009.

[2] B. An, V. Lesser, and K. M. Sim. Decommitment in multi-resource negotiation. In *Proc. of the International Joint Conference on Autonomous Agents and Multi-Agent Systems(AAMAS)*, pages 1553–1556, 2008.

[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] T. D. Nguyen and N. R. Jennings. Coordinating multiple concurrent negotiations. In *Proc. of the 3rd Int. Conf. on Autonomous Agents and Multi-Agent Systems*, pages 1064–1071, 2004.

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

[6] B. von Stengel. *Algorithmic Game Theory*, chapter Equilibrium Computation for Two-Player Games in Strategic and Extensive Form. Cambridge Univ. Press, 2007.

1. Note that in equilibrium, when a buyer $\mathbf{b}_j$ with q^{th} highest reserve price is accepting an offer with price x, the number of sellers proposing x at $t-1$ should be no less than q. The proof is omitted as it can be easily derived from the process of calculating agents' optimal prices.