

Optimal Negotiation Strategies for Agents with Incomplete Information

S. Shaheen Fatima¹ Michael Wooldridge¹ Nicholas R. Jennings²

¹ Department of Computer Science,
University of Liverpool, Liverpool L69 7ZF, U.K.
{S.S.Fatima, M.J.Wooldridge}@csc.liv.ac.uk

² Department of Electronics and Computer Science,
University of Southampton, Southampton SO17 1BJ, U.K.
nrj@ecs.soton.ac.uk

Abstract. This paper analyzes the process of automated negotiation between two competitive agents that have firm deadlines and incomplete information about their opponent. Generally speaking, the outcome of a negotiation depends on many parameters—including the agents' preferences, their reservation limits, their attitude toward time and the strategies they use. Although in most realistic situations it is not possible for agents to have complete information about each of these parameters for its opponent, it is not uncommon for agents to have partial information about some of them. Under such uncertainty, our aim is to determine how an agent can exploit its available information to select an *optimal strategy*. Here, in particular, the optimal strategies are determined considering all possible ways in which *time* can effect negotiation. Moreover, we list the conditions for convergence when both agents use their respective optimal strategies and study the effect of time on negotiation outcome.

1 Introduction

Automated negotiation is a key form of interaction in systems composed of autonomous agents [3]. Given its ubiquity, such negotiations exist in many different shapes and forms (see [7] for a taxonomy). Here, however, we consider a particular class of automated negotiation; namely, competitive bargaining over a single issue (price) between two agents that both have firm deadlines. This is exemplified by the e-commerce scenario in which a buyer agent and a seller agent negotiate over the price of a good or service. The buyer clearly prefers a low price, while the seller prefers a high one (hence the competitive nature of the encounter). In addition to attempting to obtain the best price, agents also usually need to ensure that negotiation ends before a certain deadline. However, the end point may not be the only way in which time influences negotiation behaviour. Consider the case in which the service is provided immediately after negotiation ends successfully (say at price P and time T). In some situations, it is not sufficient merely for an agent to ensure that T is any time less than its deadline. This may be the case, for instance, because one of the agents, say the buyer, could be losing utility with time as a result of not getting the service. On the other hand, the seller may perhaps gain more utility by providing the service as late as possible. Thus, in this case, the seller

tries to maximize T (within the limit of its deadline) and the buyer tries to minimize T . In short, it is clear that agents can have different attitudes toward time.

Generally speaking, the most common time effects in bargaining situations are [6]:

- Discounting: Benefits received immediately by an agent are preferred to the same benefits received in the future.
- Bargaining Cost: The bargaining process itself may incur some cost to an agent.
- Sudden Termination: An agent may have a deadline beyond which it cannot continue negotiation.

In addition to time, the outcome of a negotiation typically depends on many other parameters; such as the agents' preferences, their reservation limits, and the strategies they use. Although in most realistic cases it is not possible for agents to have complete information about all of these parameters for its opponent, it is not uncommon to have partial information about some of them. For instance, an agent may have information about its opponent's preferences, or its deadline. In this paper, we focus on situations where an agent has the following information about its opponent ¹:

- A set of possible values for the opponent's reservation limit and a binary probability distribution over these values.
- A set of possible values for the opponent's deadline and a binary probability distribution over these values.

With this information an agent can optimize its utility from price and time. However we do not assume that agents have full information about the preferences of their opponent or the strategy that they use. It is known (common knowledge) that both agents use a strategy that varies their negotiation stance with time, but the particular type of *time dependent strategy* that an agent uses is not known to its opponent. Under such uncertainty, our aim is to determine how an agent can exploit the available information to select a strategy that maximises its expected utility. Moreover, when both agents have this information about one another, we determine the impact of this information on the outcome of negotiation. This analysis is important because it enables us to construct software agents that will *optimally* negotiate on behalf of users given their state of knowledge in a given context.

The remainder of the paper is structured in the following manner. Section 2 discusses related work. Section 3 describes the basics of our negotiation model. Section 4 determines the optimal strategies for agents with incomplete information about each other. In section 5 we analyze the outcome of negotiation when both agents use their respective optimal strategies. Finally, in section 6 we present the conclusions and outline some avenues for future work.

2 Related Work

Game theoretic research typically deals with coordination and negotiation issues by assuming that agents have complete information about each other and then giving

¹ This information is *private* in the sense that the values that an agent has about its opponent are not known to the opponent.

pre-computed solutions to specific problems [10]. However this perfect information assumption is limiting because uncertainty is endemic in most realistic applications. Harsanyi et al [2] give a generalized solution for two person bargaining games with incomplete information. However there is no notion of timing issues in their model. Another important model of strategic bargaining is the infinite horizon alternating offer game [11]. Since this has a unique solution, where agents agree on a split immediately, it has been applied to automated negotiation [5]. However while this model takes time into consideration, it again assumes perfect information. Faratin et al.'s negotiation framework [1] models time as agents' deadlines and is not based on the assumption that agents have perfect information. However in this model the agent's utility functions depend only on negotiation issues like price and quality, but are independent of time.

Perhaps the work that is most closely related to ours is that of Sandholm and Vulkan [12]. In their work on bargaining with deadlines, they consider the probability distribution over agent deadlines to be common knowledge. Specifically, they address the problem of splitting the price-surplus which is known to both agents. They show that the optimal strategy is one in which agents wait until the first deadline, at which point one agent concedes everything to the other. Thus agents only ever make two offers, they either demand the entire surplus or no surplus at all. This gives the entire surplus of price to the agent with the longer deadline. But because of the offers made by agents, the deadline effect completely overrides time discounting: an agent's payoff does not change with its discounting factor. In contrast, in our work we take a binary probability distribution over agent deadlines, but we consider this to be private knowledge. In addition to this, we also take a binary probability distribution over the price-surplus. But the agents do not know their opponent's bargaining cost or discounting factor. Our optimal strategies too give the entire surplus of price to the agent with the longer deadline. However because of the difference in initial offers, our results bring out the difference between the effect of deadlines and time discounting. That is, the deadline effect on payoffs to agents does not suppress the effect of time discounting.

3 The Negotiation Model

We use an alternating offers negotiation protocol for our study. Let b denote the buyer, s the seller and let $[P_{min}^a, P_{max}^a]$ denote the range of values for price that are acceptable to agent a , where $a \in \{b, s\}$. A value for price that is acceptable to both b and s , lies between their reservation prices, i.e., in the interval $[P_{min}^s, P_{max}^b]$. The difference between P_{max}^b and P_{min}^s is the *price-surplus*. Let $p_{b \rightarrow s}^t$ denote the price offered by agent b to agent s at time t . Negotiation starts when the first offer is made by an agent. When an agent, say s , receives an offer from agent b at time t , i.e., $p_{b \rightarrow s}^t$, it rates the offer using its utility function U^s . If the value of U^s for $p_{b \rightarrow s}^t$ at time t is greater than the value of the counter-offer agent s is ready to send at time t' , i.e., $p_{s \rightarrow b}^{t'}$ with $t' > t$, then agent s accepts. Otherwise a counter-offer is made. Thus the action A that agent s takes at time t is defined as:

$$A^s(t', p_{b \rightarrow s}^t) = \begin{cases} Quit & \text{if } t > T^s \text{ where } T^s \text{ is the seller's deadline} \\ Accept & \text{if } U^s \text{ from } p_{b \rightarrow s}^t \geq U^s \text{ from counter-offer } p_{s \rightarrow b}^{t'} \\ p_{s \rightarrow b}^{t'} & \text{otherwise} \end{cases}$$




Fig. 1. Functions for the computation of $F(t)$

Given the fact that both agents have (different) deadlines, we assume that both agents use a strategy that varies their negotiation behaviour with respect to the passage of time. Thus, time is the predominant factor used to decide which value to offer in the next negotiation move. Here such strategies are called time-dependent (from [1]). Before determining the optimal strategies (in section 4), we briefly introduce the different types of time dependent strategies that we consider. These strategies vary the value of price depending on the remaining negotiation time, modeled as the above defined constant T^a . As in [1], the initial offer is a point in the interval $[P_{min}^a, P_{max}^a]$. Agents define a constant k^a that multiplied by the size of interval determines the price to be offered in the first proposal by agent a . Also the offer made by agent a to agent b at time t ($0 < t \leq T^a$) is then modeled as a function F^a depending on time as follows:

$$p_{a \rightarrow b}^t = \begin{cases} P_{min}^a + F^a(t)(P_{max}^a - P_{min}^a) & \text{for the buyer} \\ P_{min}^a + (1 - F^a(t))(P_{max}^a - P_{min}^a) & \text{for the seller.} \end{cases}$$

A wide range of time dependent functions can be defined by varying the way in which $F^a(t)$ is computed. However, functions must ensure that $0 \leq F^a(t) \leq 1$, $F^a(0) = k^a$ and $F^a(T^a) = 1$. That is, the offer will always be between the value range, at the beginning it will give the initial constant and when the deadline is reached it will offer the reservation value. Function $F^a(t)$ is defined as follows [1]:

$$F^a(t) = k^a + (1 - k^a) \left(\frac{\min(t, T^a)}{T^a} \right)^{\frac{1}{e}}$$

These families of functions represent an infinite number of possible strategies, one for each value of e . However, depending on the value of e , two extreme sets show clearly different patterns of behaviour (see Fig. 1):

1. *Boulware* [9]: For this strategy $e < 1$ and the initial offer is maintained till time is almost exhausted, when the agent concedes up to its reservation value.

2. *Conceder* [8]: For this strategy $e > 1$ and the agent goes to its reservation value very quickly. When $e = 1$ the price is increased linearly.

The value of a counter-offer depends on the initial price (IP) at which the agent starts negotiation, the final price (FP) beyond which the agent does not concede, e and T^a . Let V be a vector of these four variables, i.e., $V = [IP, FP, e, T^a]$. We call this the *counter-offer vector*. Let P^b denote P_{max}^b , P^s denote P_{min}^s , $P \in [P^b, P^s]$ and $T \in [0, T^a]$. The *negotiation outcome* O is an element of $\{(P, T), C\}$, where the pair (P, T) denotes the *price* and *time* at which agreement is reached and C denotes the conflict situation.

Agents' utility functions. The utility derived by agents depends on the final agreement on the price P and the duration of negotiation T . However, utility from price to an agent is independent of its utility from time, i.e., the buyer always prefers a low price and the seller always prefers a high price. Thus:

$$U^a : P \times T \rightarrow \Re \quad a \in (b, s)$$

We consider the following two von Neumann-Morgenstern utility functions [4] as they incorporate the effects of discounting and bargaining costs:

1. *Additive form:*

$$U^a(P, T) = U_p^a(P) + U_t^a(T)$$

where U_p^a and U_t^a are unidimensional utility functions. This form is adequate when the only effect of time is a bargaining cost which is independent of the final outcome. We defined U_p^a as $U_p^b(P) = (P^b - P)$ for the buyer and $U_p^s(P) = (P - P^s)$ for the seller. U_t^a was defined as $U_t^a(T) = c^a T$. Thus when ($c^a > 0$) the agent gains utility with time and when ($c^a < 0$) the agent loses utility with time.

2. *Multiplicative form:*

$$U^a(P, T) = U_p^a(P)U_t^a(T)$$

where, as before, U_p^a and U_t^a are unidimensional utility functions. Here preferences for attribute P , given the other attribute T , do not depend on the level of T . This form is adequate when the effects of time are bargaining cost and discounting. U_p^a was defined as before and U_t^a was defined as $U_t^a(T) = (c^a)^T$. Thus when ($c^a > 1$) the agent gains utility with time and when ($c^a < 1$) the agent loses utility with time.

Agent a 's utility from conflict is defined as $U^a(C) = 0$.

4 Optimal Negotiation Strategies

An agent's negotiation strategy defines the sequence of actions it takes during the course of negotiation. In our case, this equates to determining the value of a counter-offer which, in turn, depends on the counter-offer vector V . The information that an agent has about the negotiation parameters is called its *negotiation environment*. In order to

determine an optimal strategy an agent needs to find values for V , on the basis of its negotiation environment, that maximize its utility. Since an agent's utility depends on two parameters, price and time (see section 3), it determines the optimal price P_o and the optimal time T_o for reaching an agreement. An optimal strategy thus makes counter-offers that result in the negotiation outcome (P_o, T_o) .

4.1 Negotiation Environments

We model the negotiation environment E^b for an agent b , as a 9-tuple

$$\langle \mathcal{T}^s, \alpha^s, \alpha_c^s, \mathcal{P}^s, \beta^s, \beta_c^s, T^b, P^b, U^b \rangle$$

where:

- \mathcal{T}^s denotes a two element vector that contains possible values for s 's deadline such that the first element T_1^s is less than the second element T_2^s ;
- α^s denotes the probability that s 's deadline is T_1^s — the probability that it is T_2^s is therefore $(1 - \alpha^s)$;
- α_c^s denotes the value of α^s on the basis of which the buyer selects an optimal strategy;
- \mathcal{P}^s denotes a two element vector that contains possible values for s 's reservation price such that the first element P_1^s is less than the second element P_2^s ;
- β^s denotes the probability that s 's reservation price is P_1^s — the probability that it is P_2^s is therefore $(1 - \beta^s)$;
- β_c^s denotes the value of β^s on the basis of which the seller selects an optimal strategy;
- T^b denotes the buyer's deadline;
- P^b denotes the buyer's reservation price;
- U^b denotes the buyer's utility which is a function of price and time that decreases with price and either increases or decreases with time.

Elements $\mathcal{T}^s, \alpha^s, \mathcal{P}^s$ and β^s represent the information that b has about s . This is the agent's private information. Our aim is to determine optimal strategies for agents by considering all possible ways in which time can effect negotiation. Thus depending on the type of utility function U^b and the relation between T_1^s, T_2^s and T^b , the following six environments are defined:

- E_1^b : When $T_2^s < T^b$ and b gains utility with time, i.e., U^b is an increasing function of time.
- E_2^b : When $T_1^s < T^b < T_2^s$ and b gains utility with time.
- E_3^b : When $T_1^s > T^b$ and b gains utility with time.
- E_4^b : When $T_2^s < T^b$ and b loses utility with time, i.e., U^b is a decreasing function of time.
- E_5^b : When $T_1^s < T^b < T_2^s$ and b loses utility with time.
- E_6^b : When $T_1^s > T^b$ and b loses utility with time.

The environment for the seller can be defined analogously.




Fig. 2. Possible buyer strategies in environments E_1 and E_2

4.2 Negotiation Strategies

We formally define an agent's strategy (S^a) as a function that maps its negotiation environment, E^a , and time, T , to the counter-offer vector at time $(T + 1)$. Thus $S^a : (E^a, T) \rightarrow V^{T+1}$. Let O be the outcome that results from strategy S^a . S^a is the optimal strategy for agent a if it maximizes $U^a(O)$.

For any environment E_i^a , S_i^a denotes the corresponding optimal strategy. The optimal strategies are determined for each of the above six environments. An optimal strategy should result in agreement at the optimal price and at the optimal time. An agent's preference for price is independent of time, i.e., the buyer always prefers a low price and the seller always prefers a high price. We therefore determine the optimal strategy assuming $\beta^s = 1$ (section 4.3). This gives the optimal time for reaching agreement. Then for this optimal time we find the strategy that gives the optimal price by taking the actual value of β^s . This is explained in section 4.4. The resulting strategy is therefore optimal in both time and price.

4.3 Optimal Strategies in Particular Environments when $\beta^s = 1$

This section details the optimal strategy for each of the environments noted above. The analysis is from the perspective of the buyer, although strategies for the seller can be defined analogously.

Environment E_1^b . In this environment both the deadlines of s are less than T^b . As both agents use time dependent tactics, s will concede up to P_1^s latest by T_1^s or T_2^s . Hence it is enough if b uses some strategy that concedes only up to P_1^s . Since utility increases with time until T^b , b gets maximum utility if agreement is reached as late as possible. Since there are two possible values for the seller's deadline, negotiation could end either at T_1^s or T_2^s and there are two possible strategies b can use in this case (see Fig. 2A). It could use a strategy (s_1^b) that starts at P_{min}^b and approaches P_1^s at T_1^s and then makes no further concessions till T_2^s , or a strategy (s_2^b) that starts at P_{min}^b and approaches P_1^s at time T_2^s . These can be defined more formally as follows:

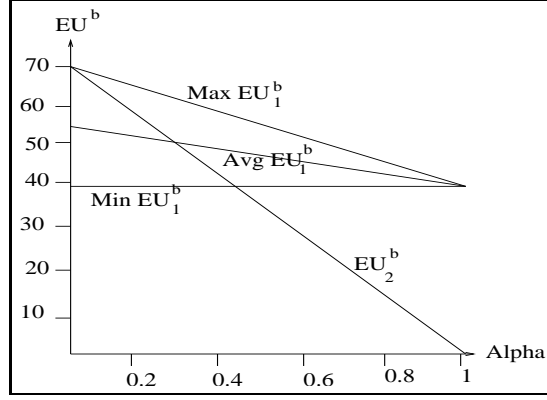


Fig. 3. Effect of α on the EU to buyer in E_1

$$s_1^b(E_1^b, T) = [P_{min}^b, P_1^s, \text{Boulware}, T_1^s] \text{ if } T \leq T_1^s \text{ and } P_1^s \text{ otherwise and}$$

$$s_2^b(E_1^b, T) = [P_{min}^b, P_1^s, \text{Boulware}, T_2^s] \text{ for all values of } T.$$

By using strategy s_1^b , b gets a utility of $U^b(P_1^s, T_1^s)$ with probability α^s and a utility of $U^b(P_1^s, T)$ with probability $(1 - \alpha^s)$. Note that T lies between T_1^s and T_2^s depending on s 's strategy. So the expected utility to b , EU^b , from strategy s_1^b is

$$EU_1^b = \alpha^s U^b(P_1^s, T_1^s) + (1 - \alpha^s) U^b(P_1^s, T)$$

where $T_1^s < T \leq T_2^s$. On the other hand, if b uses strategy s_2^b , it gets a utility of $U^b(C)$ with probability α^s and a utility of $U_b(P_1^s, T_2^s)$ with probability $(1 - \alpha^s)$. So its expected utility from strategy s_2^b becomes

$$EU_2^b = \alpha^s U^b(C) + (1 - \alpha^s) U^b(P_1^s, T_2^s).$$

Turning now to an experimental evaluation of these strategies. Agent b 's utility was defined as follows ²

$$U^b(P, T) = \begin{cases} \min(1, \frac{P^b - P}{P^b - P_{min}^b})T & \text{for } T \leq T^b \\ 0 & \text{for } T > T^b \end{cases}$$

Thus b 's utility linearly increases with price and time where $\min(1, \frac{P^b - P}{P^b - P_{min}^b})$ gives a value between 0 and 1. The expected utility to b from any strategy depends on the following parameters: $[T_1^s, T_2^s, T^b, P_1^s, P^b, P_{min}^b, \alpha^s]$. Let θ_t^s denote the length of time interval between T_1^s and T_2^s . For all contexts we assigned $T_1^s = 30$, $P^b = 70$, $P_{min}^b = 20$, $P_1^s = 30$ and $T^b = 100$. In our experiments we computed the expected utility for

² In all the experiments reported in this section we obtained similar results for additive and multiplicative forms of utility functions. We therefore describe here the experiments for only one multiplicative form.




Fig. 4. Effect of α on the EU to buyer in E_2

different values of θ_t^s . For each value of θ_t^s we varied α^s between 0 and 1 and found the EU^b in each case. In the first term for EU^b from strategy s_1^b , both parameters to U^b , i.e., P_1^s and T_1^s , are constants. In the second term, P_1^s is constant but T could be any value in the interval T_1^s to T_2^s depending on the strategy that the seller uses. As a result we get a range of values for EU^b from strategy s_1^b . Since utility to b is an increasing function of time, for the same price P_1^s , b gets maximum utility at T_2^s and minimum utility at T_1^s . All other possible values for EU^b lie within this range. We therefore computed the average EU^b for every value of α^s . All the parameters to U^b in the EU^b from strategy s_2^b are constants. This gives a single value for EU^b for every value of α^s .

The results of this experiment for $\theta_t^s = 40$ are as shown in Fig. 3. For values of α^s up to 0.3, strategy s_2^b gives a higher EU^b than the average EU^b from strategy s_1^b . For higher values of α^s , the average EU^b from s_1^b becomes higher than the EU^b from s_2^b . Let α_c^s denote the value of α^s below which s_2^b is better than s_1^b and above which s_1^b is better than s_2^b . This value of α_c^s depends on the value of θ_t^s . As we increased θ_t^s , there was a corresponding increase in the value of α_c^s . A decrease in θ_t^s resulted in a corresponding decrease in α_c^s . Thus the optimal strategy S_1^b is s_2^b when $\alpha^s < \alpha_c^s$ and s_1^b when $\alpha^s > \alpha_c^s$.

Environment E_2^b . As one of the possible deadlines for the seller, T_2^s , is greater than T^b , b needs to concede up to P^b in order to maximize its chances of reaching an agreement. Again it could use one of two possible strategies (see Fig. 2B). A strategy s_1^b that starts at P_{min}^b and approaches P^b at T^b by crossing P_1^s at T_1^s , or a strategy s_2^b that starts at P_{min}^b and approaches P^b at T^b , crossing P_1^s somewhere in the interval between T_1^s and T^b . That is,

$$s_1^b(E_2^b, T) = \begin{cases} [P_{min}^b, P_1^s, \text{Boulware}, T_1^s] & \text{if } T \leq T_1^s \\ [P_1^s, P^b, \text{Boulware}, (T^b - T_1^s)] & \text{if } T > T_1^s \end{cases}$$

$$s_2^b(E_2^b, T) = [P_{min}^b, P^b, \text{Boulware}, T^b] \text{ for all values of } T.$$

By using s_1^b , b gets a utility of $U^b(P_1^s, T_1^s)$ with probability α^s and utility $[\lambda U^b(P, T) +$




Fig. 5. Optimal buyer strategies when utility decreases with time

$(1-\lambda)U^b(C)]$ with probability $(1-\alpha^s)$. λ denotes the probability that the seller crosses P^b prior to T^b . So the expected utility to b from strategy s_1^b is

$$EU_1^b = \alpha^s U^b(P_1^s, T_1^s) + (1 - \alpha^s)[\lambda U^b(P, T) + (1 - \lambda)U^b(C)]$$

where $(P_1^s < P < P^b)$ and $(T_1^s < T < T^b)$.

If b uses strategy s_2^b it gets a utility of $U^b(C)$ with probability α^s and a utility of $[\lambda U^b(P, T) + (1 - \lambda)U^b(C)]$ with probability $(1 - \alpha^s)$ where $(P_1^s < P < P^b)$ and $(T_1^s < T < T^b)$. So b 's expected utility from s_2^b becomes

$$EU_2^b = \alpha^s U^b(C) + (1 - \alpha^s)[\lambda U^b(P, T) + (1 - \lambda)U^b(C)]$$

In our experimental evaluation of these strategies, we get a range of values for EU^b from both strategies s_1^b and s_2^b . We assigned $\lambda = 1$. The results of our experiments (for the same values of all other variables as in environment 1) are shown in Fig. 4. The average EU^b from strategy s_1^b was higher than the average EU^b from s_2^b for all values of α^s . Note that in this case the relationship between s_1^b and s_2^b does not depend on the value of θ_t^s . Thus the optimal strategy S_2^b is s_1^b for all values of α^s .

Environment E_3^b . Since both possible deadlines for the seller are greater than T^b , b has to use a strategy S_3^b that starts at P_{min}^b and concedes up to P^b by T^b . Thus for all values of T

$$S_3^b(E_3^b, T) = [P_{min}^b, P^b, Boulware, T^b]$$

In the remaining three environments b 's utility decreases with time, and it can maximize its utility by using a strategy that ends negotiation as early as possible.

Environment E_4^b . Here $T_2^s < T^b$ and the buyer maximizes its utility by minimizing price by conceding only up to P_1^s since it knows from its environment that s will offer P_1^s latest by T_1^s or T_2^s . In addition to this it can maximize its utility from time by offering




Fig. 6. Possible buyer strategies when $\beta \neq 1$

P_1^s as early as possible. Thus b 's best strategy S_4^b is to offer P_1^s every time starting from its first offer (see Fig. 5A). Thus

$$S_4^b(E_4^b, T) = [P_{min}^b, P_1^s, Conceder, T_2^s]$$

for all values of T and the counter-offers are always P_1^s .

Environment E_5^b . In this environment b 's best strategy is to offer P_1^s from the start of negotiation till T_1^s . If T_1^s is the actual deadline for the seller, negotiation would end latest by T_1^s at price P_1^s . On the other hand, if T_2^s is the seller's actual deadline, negotiation could continue beyond T_1^s (depending on the seller's strategy). b now has to concede up to P^b at the end of T_1^s since T^b is less than T_2^s . The strategy S_5^b is shown in Fig. 5B. Thus

$$S_5^b(E_5^b, T) = \begin{cases} [P_{min}^b, P_1^s, Conceder, T_1^s] & \text{if } T \leq T_1^s \\ [P_1^s, P^b, Conceder, (T^b - T_1^s)] & \text{if } T > T_1^s \end{cases}$$

Environment E_6^b . In this environment b 's best strategy S_6^b is to concede up to P^b at the beginning of negotiation since $U^b(P^b, 1) \geq U^b(P^s, 2)$. Thus

$$S_6^b(E_6^b, T) = [P_{min}^b, P^b, Conceder, T^b]$$

for all values of T and the counter-offers are P^b throughout negotiation.

4.4 Optimal Strategies in Particular Environments when $\beta^s \neq 1$

The above strategies give an optimal value for T if $\beta^s = 1$. When $\beta^s \neq 1$ an optimal value for price needs to be determined. There are two possible values for the seller's reservation limit, P_1^s with probability β^s and P_2^s with probability $(1 - \beta^s)$. b now has two possible strategies (see Fig. 6). It can concede up to P_1^s by T (strategy S_1^b) and get an EU^b of

$$EU_1^b = \beta^s U^b(P_1^s, T) + (1 - \beta^s) U^b(C)$$

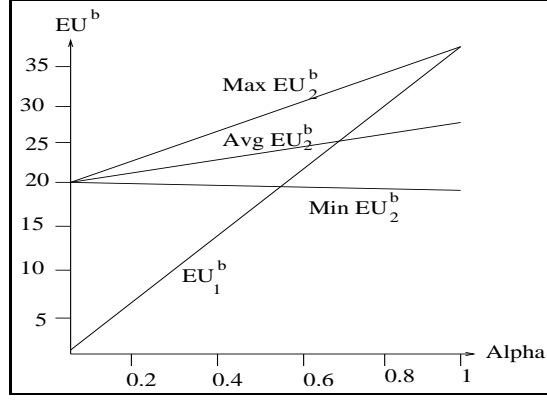


Fig. 7. Effect of β on the EU to buyer

or concede up to P_2^s by T (strategy s_2^b) and get an EU^b of

$$EU_2^b = \beta^s U^b(P, T) + (1 - \beta^s) U^b(P_2^s, T)$$

where $(P_1^s < P < P_2^s)$.

In terms of our experimental evaluation, the EU^b from any strategy depends on the following parameters: $[P_1^s, P_2^s, P^b, P_{min}^b, T, \beta^s]$. We assigned $P_1^s = 30, P_2^s = 50$. For all contexts we assigned $P^b = 70, P_{min}^b = 20, T = 50$ and $T^b = 100$. We varied β^s between 0 and 1 and computed the EU^b for both strategies for each value of β^s . The results of this experiment are as shown in Fig. 7. For values of β^s up to 0.65 the average utility from s_2^b was higher than the EU^b from strategy s_1^b . For values of β^s higher than 0.65, the EU^b from strategy s_1^b became higher than the average EU^b from strategy s_2^b . Let β_c^s denote the value of β^s below which s_2^b is better than s_1^b and above which s_1^b is better than s_2^b . Let θ_s^p denote the difference between P_1^s and P_2^s . The value of β_c^s depends on the value of θ_s^p . As we increased θ_s^p there was a corresponding decrease in the value of β_c^s . The decrease in θ_s^p resulted in a corresponding increase in β_c^s . Thus for $(\beta^s < \beta_c^s)$ the optimal price is P_2^s and for $(\beta^s > \beta_c^s)$ the optimal price is P_1^s .

5 Negotiation Outcomes

We now determine the outcome of negotiation when both agents use their respective optimal strategies as determined in table 1. Since the buyer and seller could either gain or lose utility with time, we have the following four possibilities:

1. Both buyer and seller gain utility with time.
2. Buyer gains and seller loses utility with time.
3. Buyer loses and seller gains utility with time.
4. Both buyer and seller lose utility with time.

The ordering on the deadlines of b and s will be one of the following:

Environment	Optimal Strategy
E_1^b	$[P_{min}^b, P_o^s, \text{Boulware}, T_2^s]$ for all values of T if $(\alpha^s < \alpha_c^s)$ $[P_{min}^b, P_o^s, \text{Boulware}, T_1^s]$ for all values of T if $(\alpha^s > \alpha_c^s)$
E_2^b	$[P_{min}^b, P_o^s, \text{Boulware}, T_1^s]$ if $T \leq T_1^s$ $[P_o^s, P^b, \text{Boulware}, (T^b - T_1^s)]$ if $T > T_1^s$
E_3^b	$[P_{min}^b, P^b, \text{Boulware}, T^b]$ for all values of T
E_4^b	$[P_{min}^b, P_o^s, \text{Conceder}, T_2^s]$ for all values of T
E_5^b	$[P_{min}^b, P_o^s, \text{Conceder}, T_1^s]$ if $T \leq T_1^s$ $[P_o^s, P^b, \text{Conceder}, (T^b - T_1^s)]$ if $T > T_1^s$
E_6^b	$[P_{min}^b, P^b, \text{Conceder}, T^b]$ for all values of T

Table 1. Optimal Strategies for the Buyer (P_o^s denotes the optimal value for the seller's reservation price)

1. $T_1^s < T_2^s < T_1^b < T_2^b$ (D1)
2. $T_1^b < T_2^b < T_1^s < T_2^s$ (D2)
3. $T_1^s < T_1^b < T_2^s < T_2^b$ (D3)
4. $T_1^s < T_1^b < T_2^b < T_2^s$ (D4)
5. $T_1^b < T_1^s < T_2^s < T_2^b$ (D5)
6. $T_1^b < T_1^s < T_2^b < T_2^s$ (D6)

We consider the case where b gains with time and s loses with time and determine the outcome of negotiation for all possible orderings of deadlines. For ordering D1, let T_1^s be s 's actual deadline (see Fig. 8A). This corresponds to E_6^s for the seller and its optimal strategy is S_6^s . For agent b it corresponds to E_1^b (with a high value for α^s) irrespective of whether its deadline is T_1^b or T_2^b . Let P_o^s and P_o^b denote the optimal values for the seller's minimum reservation price and buyer's maximum reservation price respectively. Thus when b uses strategy S_1^b , s uses S_6^s , and the outcome of negotiation becomes (P_o^s, T_1^s) . On the other hand, if T_2^s is s 's actual deadline (see Fig. 8B), it corresponds to E_6^s for the seller and its optimal strategy is S_6^s . For the buyer it is E_1^b and its optimal strategy is S_1^b (α^s is low). The negotiation outcome now becomes (P_o^s, T_2^s) . In both cases, since b 's deadline is greater than s 's deadline, it gets the entire surplus of price.

D2 is similar to D1 except that b and s are interchanged and the negotiation outcome is (P_o^b, T_1^b) when T_1^b is b 's actual deadline and (P_o^b, T_2^b) when T_2^b is b 's actual deadline (see Fig. 8C and 8D).

For D3, let T_1^s be s 's actual deadline (see Fig. 8E). This corresponds to E_6^s for s and its optimal strategy is S_6^s . For the buyer it is E_2^b if its deadline is T_1^b and E_1^b (with α^s high) if its deadline is T_2^b . Using the corresponding optimal strategies, in both cases the negotiation outcome is (P_o^s, T_1^s) . If T_2^s is s 's actual deadline (see Fig. 8F), this corresponds to E_5^s for s and its optimal strategy is S_5^s . Now if T_1^b is b 's actual deadline (E_2^b), its optimal strategy is S_2^b and the negotiation outcome is (P_o^b, T_1^b) ; the entire surplus of price goes to S . But if T_2^b is b 's actual deadline (E_1^b with α^s low), its optimal strategy gives the outcome (P_o^s, T_2^s) and the entire surplus of price goes to b .

For D4, let T_1^s be the seller's actual deadline (see Fig. 8 G). This corresponds to E_6^s for the seller and its optimal strategy is S_6^s . If T_1^b or T_2^b is b 's actual deadline (E_2^b),




Fig. 8. Negotiation Outcome (thick lines denote seller strategy and dotted lines buyer strategy)

then its optimal strategy is S_2^b in both cases. The negotiation outcome is (P_o^s, T_1^s) . The entire surplus of price goes to b . On the other hand, if T_2^s is s 's actual deadline (E_4^b), its optimal strategy is S_4^b (see Fig. 8H). The optimal strategy for b is S_2^b irrespective of whether its actual deadline is T_1^b or T_2^b . The corresponding outcome of negotiation is (P_o^b, T_1^b) or (P_o^b, T_2^b) , and the entire surplus goes to s . In the same way the outcome of negotiation, if both agents use the optimal strategies as determined in section 3, can be found for all the remaining cases. These results are summarized in table 2. T_0 denotes the beginning of negotiation.

As seen from table 2, the negotiation outcome remains the same in the first three cases, i.e., when both agents gain utility with time, or when any one of them gains and the other loses with time. This happens because the agent that gains utility with time delays in making an offer that is acceptable to the opponent till the earliest deadline is reached. When both agents lose utility with time, they make the maximum required concession at the earliest opportunity, i.e. the start of negotiation. Agreement is therefore reached at the beginning of negotiation.

5.1 Conditions for convergence of optimal strategies

The values of α_c^s and β_c^s are crucial in determining b 's optimal strategy. For instance in negotiation environment E_1^b , the agent's optimal strategy S_1^b is s_2^b when $(\alpha^s < \alpha_c^s)$ and S_1^b is s_1^b when $(\alpha^s > \alpha_c^s)$. Thus in order for the strategies to converge, an agent, say b 's, degree of belief α^s about s 's deadline must be less than α_c^s when the actual deadline for s is T_2^s and it must be greater than α_c^s when the actual deadline for s is T_1^s .

Deadline Ordering	Seller's Deadline	Case 1, 2 and 3		Case 4	
		Negotiation Outcome		Negotiation Outcome	
		Buyer's Deadline		Buyer's Deadline	
D1	T_1^s	(P_o^s, T_1^s)	(P_o^s, T_1^s)	(P_o^s, T_0)	(P_o^s, T_0)
		T_1^b	T_2^b	T_1^b	T_2^b
	T_2^s	(P_o^s, T_2^s)	(P_o^s, T_2^s)	(P_o^s, T_0)	(P_o^s, T_0)
		T_1^b	T_2^b	T_1^b	T_2^b
D2	T_1^s	(P_o^b, T_1^b)	(P_o^b, T_2^b)	(P_o^b, T_0)	(P_o^b, T_0)
		T_1^b	T_2^b	T_1^b	T_2^b
	T_2^s	(P_o^b, T_1^b)	(P_o^b, T_2^b)	(P_o^b, T_0)	(P_o^b, T_0)
		T_1^b	T_2^b	T_1^b	T_2^b
D3	T_1^s	(P_o^s, T_1^s)	(P_o^s, T_1^s)	(P_o^s, T_0)	(P_o^s, T_0)
		T_1^b	T_2^b	T_1^b	T_2^b
	T_2^s	(P_o^s, T_1^b)	(P_o^s, T_2^s)	(P_o^s, T_0)	(P_o^s, T_0)
		T_1^b	T_2^b	T_1^b	T_2^b
D4	T_1^s	(P_o^s, T_1^s)	(P_o^s, T_1^s)	(P_o^s, T_0)	(P_o^s, T_0)
		T_1^b	T_2^b	T_1^b	T_2^b
	T_2^s	(P_o^b, T_1^b)	(P_o^b, T_2^b)	(P_o^b, T_0)	(P_o^b, T_0)
		T_1^b	T_2^b	T_1^b	T_2^b
D5	T_1^s	(P_o^b, T_1^b)	(P_o^s, T_1^s)	(P_o^b, T_0)	(P_o^s, T_0)
		T_1^b	T_2^b	T_1^b	T_2^b
	T_2^s	(P_o^b, T_1^b)	(P_o^s, T_2^s)	(P_o^b, T_0)	(P_o^s, T_0)
		T_1^b	T_2^b	T_1^b	T_2^b
D6	T_1^s	(P_o^b, T_1^b)	(P_o^s, T_1^s)	(P_o^b, T_0)	(P_o^s, T_0)
		T_1^b	T_2^b	T_1^b	T_2^b
	T_2^s	(P_o^b, T_1^b)	(P_o^b, T_2^b)	(P_o^b, T_0)	(P_o^b, T_0)
		T_1^b	T_2^b	T_1^b	T_2^b

Table 2. Outcome of negotiation when both agents use their respective optimal strategies

Similarly the optimal value for price depends on the value of β_c^s . For instance, for the buyer, the optimal price is P_2^s if $(\beta^s < \beta_c^s)$ and it is P_1^s if $(\beta^s > \beta_c^s)$. Thus convergence is guaranteed only when the values of α^s and β^s satisfy the following conditions:

1. $(\alpha^s < \alpha_c^s)$ if $(T^s = T_2^s)$ and $(\alpha^s > \alpha_c^s)$ if $(T^s = T_1^s)$ for the buyer.
2. $(\alpha^b < \alpha_c^b)$ if $(T^b = T_2^b)$ and $(\alpha^b > \alpha_c^b)$ if $(T^b = T_1^b)$ for the seller.
3. $(\beta^s < \beta_c^s)$ if $(P^s = P_2^s)$ and $(\beta^s > \beta_c^s)$ if $(P^s = P_1^s)$ for the buyer.
4. $(\beta^b < \beta_c^b)$ if $(P^b = P_2^b)$ and $(\beta^b > \beta_c^b)$ if $(P^b = P_1^b)$ for the seller.

5.2 Payoffs to agents

The utility that an outcome yields to an agent is a function of price and time. The factors that determine the outcome of negotiation are the agents' deadlines and their discounting or bargaining costs. We therefore analyze their effect on the utilities to both agents.

Effect of deadlines on the payoffs. As seen from table 2, when agents have unequal deadlines, the entire surplus of price goes to the agent with the longer deadline.

Effect of discounting/bargaining costs on payoffs. When at least one agent gains utility with time, agreement is reached at the earlier deadline; when both lose on time, agreement is reached toward the beginning of negotiation. Let the pair (U_b^1, U_s^1) denote the utility to b and s when both gain with time, (U_b^2, U_s^2) when b gains and s loses with time, (U_b^3, U_s^3) when b loses and s gains with time and (U_b^4, U_s^4) when both lose with time. When the effect of time on the two agents is not identical, i.e., when one of them gains utility with time and the other loses, the following result was observed. The agent that gains on time gets the same utility as compared to the case where both gain with time, i.e., for the buyer $(U_b^1 = U_b^2)$ and for the seller $(U_s^1 = U_s^3)$. The agent that loses on time gets less utility with time gets a lower utility when compared to the case where both gain or both lose with time, i.e., for the buyer $(U_b^3 < U_b^4)$ and $(U_b^3 < U_b^1)$ and for the seller $(U_s^2 < U_s^4)$ and $(U_s^2 < U_s^1)$.

6 Conclusions

This paper determined what the optimal negotiation strategies are for agents that find themselves in environments with different information states. Specifically, we considered situations where agents have uncertain information about two negotiation parameters (the opponent's deadline and reservation limit) but do not have any information about their opponent's bargaining cost, discounting factor or strategy. We listed conditions for convergence of these optimal strategies and studied the effect of time on the negotiation outcome. In the future we intend to extend our analysis to determine if this mutual strategic behavior leads to equilibria and then analyze situations where agents have limited information about other negotiation parameters like the opponent's bargaining cost, its discounting factor or its strategy to compare their relative influences on the negotiation outcome.

Acknowledgments: This work was supported by the EPSRC under grant GR/M07076 ("Practical Negotiation for Electronic Commerce").

References

1. P. Faratin, C. Sierra, and N. R. Jennings. Negotiation decision functions for autonomous agents. *International Journal of Robotics and Autonomous Systems*, 24(3-4):159–182, 1998.
2. J.C. Harsanyi and R. Selten. A generalized nash solution for two-person bargaining games with incomplete information. *Management Science*, 18(5):80–106, January 1972.
3. N. R. Jennings, P. Faratin, A. R. Lomuscio, S. Parsons, C. Sierra, and M. Wooldridge. Automated negotiation: Prospects, methods and challenges. *Int Journal of Group Decision and Negotiation*, 10(2):199–215, 2001.
4. R. Keeney and H. Raiffa. *Decisions with Multiple Objectives: Preferences and Value Trade-offs*. New York: John Wiley, 1976.
5. S. Kraus, J. Wilkenfeld, and G. Zlotkin. Negotiation under time constraints. *Artificial Intelligence Journal*, 75(2):297–345, 1995.

6. Z. A. Livne. *The Role of Time in Negotiations*. PhD thesis, Massachusetts Institute of Technology, 1979.
7. A. Lomuscio, M. Wooldridge, and N. R. Jennings. A classification scheme for negotiation in electronic commerce. In F. Dignum and C. Sierra, editors, *Agent-Mediated Electronic Commerce: A European AgentLink Perspective*, pages 19–33. Springer Verlag, 2001.
8. D. G. Pruitt. *Negotiation Behavior*. Academic Press, 1981.
9. H. Raiffa. *The Art and Science of Negotiation*. Harvard University Press, Cambridge, USA, 1982.
10. J. S. Rosenchein and G. Zlotkin. *Rules of Encounter*. MIT Press, 1994.
11. A. Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):97–109, January 1982.
12. T. Sandholm and N. Vulkan. Bargaining with deadlines. In *AAAI-99*, pages 44–51, Orlando, FL, 1999.