

A Solution for Bilateral Negotiations in the Navy
Detailing Process

for the
Navy Detailing Process,
Cognitive Agents Technology Project

Cuihong Li
cuihong@andrew.cmu.edu

Joseph Giampapa
garof@cs.cmu.edu

Katia Sycara
katia@cs.cmu.edu

December 23, 2003

Abstract

We provide an integrative solution for negotiations in the Navy detailing process considering the uncertain and dynamic outside options. The outside options influence the negotiation strategies via the impact on the reservation prices. The solution is composed of three modular models: single-threaded negotiations, synchronized multi-threaded negotiations, and dynamic multi-threaded negotiations. The single-threaded negotiation model provides the negotiation strategy given the reservation price. The other two models calculate the reservation price of a negotiation thread based on the model of the outside options by viewing all outside options as a multi-threaded negotiation process. The model of synchronized multi-threaded negotiations considers the presence of concurrently available outside options and provides an approach to estimate the outcome when the threads are known. The model of dynamic multi-threaded negotiations expands the synchronized model by considering the uncertain outside options that may come dynamically in the future. We propose two effective negotiation strategies, the time-dependent strategy and Bayesian learning strategy, from the AI field. Four heuristic approaches are designed to estimate the expected utility from a synchronized multi-threaded negotiation. A Poisson process is used to model the random sequential arrivals, and formulas are provided to calculate the expected utility from a negotiation process when uncertain outside options may come in the future.

We give preliminary experimental analysis to characterize the impact of outside options on the reservation price and so on the negotiation strategy. The results show that the utility of a negotiator improves significantly when she considers outside options from when she does not, and the average utility is higher when she both considers the concurrent outside options and foresee the future ones than when she only considers the concurrent outside options.

1 Introduction

Bilateral negotiations is an important mechanism to implement flexible and distributed matching in the Navy detailing system [4, 7]. In the Navy detailing process a negotiator can face more than one potential matching alternative. For example, a Command may find more than one Sailor that is qualified for the job, and a Sailor can be informed of more than one job vacancies that interest him. These alternatives are called *outside options*. Accepting a proposal in one negotiation means refusing all outside options. On the other hand one may leave a negotiation (called “opt-out” of a negotiation) without reaching an agreement based on the expectation of reaching a more favorable agreement in outside options. Modelling the outside options and understanding the interaction between outside options and a negotiation process is an essential aspect to designing an effective negotiation strategy in the Navy detailing process.

Outside options can exist *concurrently* with a negotiation, or come *sequentially* in the future [6]. A *concurrently* available outside option is a negotiation thread that the negotiator is involved in simultaneously with another thread. This happens because a Command may find multiple potential Sailors that are available for negotiations for the same job at the same time. A Sailor may also be invited to more than one negotiation - one for each potential job - simultaneously. A *sequentially* available outside option is a matching opportunity that comes in the future. A Command is not informed at one time of all potential Sailors that will become available during the whole search period, neither is a Sailor aware of all interesting job vacancies during their application period. The information on Sailors and jobs that are available is published periodically and sequentially. To obtain information may also induce searching cost, which prevents awareness of complete information at one time. Outside options are *uncertain* in terms of both *availability* and *quality*. The *availability* of outside options is uncertain because a negotiator is not sure when an outside option is available and how many are available. The *quality* of outside options is uncertain because a negotiator cannot predict the outcome of a negotiation thread, or the preferences of the other party in a negotiation thread. How to model the availability and uncertainty of outside options is an important consideration for modelling.

Outside options affect the negotiation strategies via their impact on the reservation price. The *reservation price* is the worst agreement that a negotiator can accept. For example, in a buyer-seller negotiation model, the reservation price of the buyer is the highest price she is willing to pay for the negotiated good. For the seller, the reservation price is the lowest price at which he is willing to sell the good. The price at which the seller is willing to sell depends on the production cost of the seller. The price at which the buyer is willing to buy depends on the valuation of the buyer to the good.

Additionally they both depend on the availability of other buyers or sellers. Let's take the standpoint of the buyer. Similar statements can be drawn for the seller. For the buyer, if there are no outside sellers, the reservation price is equal to the valuation . However, the negotiation does not necessarily end up with the reservation price because the seller does not know the buyer's reservation price. If there are other sellers joining in the market and the original seller is not a monopoly any more, then the buyer will decrease her reservation price hoping that she could reach a deal with other sellers, if she cannot reach an agreement with the current seller with a price less than the reservation price. The reservation price will be lower if the buyer expects that there are more outside sellers with possibly lower prices. The reservation price impacts the proposal and response decisions of a negotiator. When there are outside options, design of an effective negotiation strategy can be divided into two parts: the first is to design a negotiation strategy given the reservation price and other inputs, the second is to calculate the reservation price based on the model of outside options.

In our previous modelling work [6] we have proposed a nested model for negotiations in the Navy detailing process considering the uncertain and dynamic outside options. The model is composed of three modules, single-threaded negotiations, synchronized multi-threaded negotiations, and dynamic multi-threaded negotiations. These three models embody increased sophistication and complexity. The single-threaded negotiation model provides the negotiation strategies without specifically considering outside options. The model of synchronized multi-threaded negotiations builds on the single-threaded negotiation model, and considers the presence of concurrently available outside options by calculating the reservation price based on the other existing negotiation threads. The model of dynamic multi-threaded negotiations expands the synchronized multi-threaded model by considering the uncertain outside options that may come dynamically in the future.

In this report we present specific solutions of these models. We propose two effective negotiation strategies, the time-dependent strategy and Bayesian learning strategy, from the AI field. Four heuristic approaches are designed to estimate the expected utility from a synchronized multi-threaded negotiation. A Poisson process is used to model the random sequential arrivals, and formulas are provided to calculate the expected utility from a negotiation process when uncertain outside options may come in the future. Empirical analysis is provided to characterize the impact of outside options on the reservation price and therefore on the negotiation strategy. The results show that the utility of a negotiator improves significantly when she considers outside options from when she does not, and the average utility is higher when she both considers the concurrent outside options and foresees the future ones than when she only considers the concurrent outside options.

The rest of the report is organized as follows: Section 2 presents the models and solutions. In Section 3 we describe the experimental results. Section 4 concludes.

2 The model

For the convenience of presentation we call the two agents in a bilateral negotiation a buyer and a seller, and we present the model from the buyer's perspective. The buyer prefers a lower value of the negotiated issue while the seller prefers a higher value. In the Navy detailing process we can regard a Command as a buyer and a Sailor as a seller. The roles of a buyer and a seller are interchangeable by changing the sign of the value of the negotiated issue.

There are T periods over the entire horizon of a detailing window. Let a period be denoted by t , $t = 0, \dots, T-1$. A buyer needs to reach an agreement with a seller before period T . The potential sellers may come in different time unexpectedly with different reservation prices, and the buyer can negotiate with the sellers simultaneously. The negotiation between the buyer and a seller is called a *negotiation thread*. The number of threads in period t is denoted by n_t , and the collection of threads in period t is denoted by $D_t = \{d_i\}_{i=1}^{n_t}$. The seller in the thread d_i is denoted by s_i . Based on the background information on the sellers (or sellers' products), the buyer can value the sellers (or the sellers' products) differently. For example, a Command can attach different values to having the job filled by different Sailors based on the Sailors' skills and experiences and the job's requisites. A Sailor can also have different preferences on different jobs because of the difference in the location or properties. Let the value of the seller s_i be v_i . If the buyer reaches an agreement with the seller s_i at x , then the *utility* of the buyer is $v_i - x$.

The buyer wants to reach the lowest possible agreement with a seller. But she does not know the bottom line of the seller, or what price is acceptable to the seller. If the price she insists on is low, she could get a high profit but risk losing the cooperation opportunity with the seller. On the other hand if the price she agrees on is high, she can make a deal with the seller with high probability but then the profit is low. Although the buyer does not know the reservation price of a seller, she can have some estimation of the information based on statistical aggregation of the historical data or survey work. The historical data records the agreements that were reached on the same or similar products (jobs) in the past. A negotiator can also, maybe by the help of a third party, do a survey to ask the reservation prices of a representative population. The estimation of the reservation price of a seller is characterized by a probability distribution $F(\cdot)$, where $F(x)$ denotes the probability that the reservation price of a

seller is no greater than x . This probability distribution is called the *prior belief* of the buyer. A negotiation provides a mechanism for the negotiators to exchange messages and adjust their proposals. Usually a negotiator will start with a favorable proposal, and then make sequential concession until a proposal is accepted or the negotiation deadline is reached. The negotiation strategy decides the pace of concession at each step based on the single-threaded negotiation model given the buyer's reservation price and the prior belief.

When there are outside options the decision of a negotiator is more complicated. The buyer will expect to reach a utility that is no worse than the expected utility that she could achieve from the outside options. In other words, the buyer has a threshold on the lowest utility that she should achieve from a negotiation thread based on the expectation of the outside options. The lowest utility to achieve in thread d_i is called the *reservation utility* OU_i in the thread. The reservation utility is equal to the expected utility from the outside options, which can be viewed together as a multi-threaded negotiation. Given the reservation utility OU_i , the reservation price R_i of the buyer in thread d_i can be calculated by $R_i = v_i - OU_i$. If there are no outside options, we can say that the reservation utility is zero, or the reservation price R_i is equal to the value v_i . Because of the heterogeneity among the sellers, the reservation prices in each thread may be different. If the reservation price in each thread is known, the buyer can apply the single-threaded negotiation model to make the negotiation decisions in each thread.

Calculation of the expected utility from the outside options depends on the model on the outside options, and on the approach to estimate the expected utility from a multi-threaded negotiation. In a synchronized multi-threaded negotiation model the outside options at period t for thread d_i are the other concurrently existing negotiation threads $D_t \setminus d_i$. The synchronized model maps the current outside options to the reservation utility $OU_i(D_t \setminus d_i)$ of each thread d_i , $i = 1, \dots, n_i$. The dynamic multi-threaded negotiation model also considers the outside options that may come in the future at uncertain times with uncertain values. Let the probability that a new opponent arrives in a period be p , and the probability distribution of the value of an opponent be $\Phi(\cdot)$, where $\Phi(v)$ is the probability that the value of an opponent is less than v . The dynamic multi-threaded negotiation model calculates the reservation utility $OU_i(D_t \setminus d_i | p, \Phi(\cdot))$ for each current thread d_i based on the current outside options $D_t \setminus d_i$, given the arrival probability and the probability distribution of opponents' values. The dynamic multi-threaded negotiation model can be viewed as a synchronized model with uncertain threads.

In Section 2.1 we first present the negotiation strategy solution in a single-threaded negotiation. In Section 2.2 the influence of the concurrent negotiation threads is quan-

tified in the reservation utility. In Section 2.3 the negotiation threads that may come sequentially in the future are considered additionally and the impact is integrated in the reservation utility and thus in the negotiation strategy.

2.1 Single-threaded negotiations

We describe the negotiations based on an alternating-offers negotiation protocol, because (1) it is a sequential negotiation protocol, which allows negotiators to dynamically adjust the offers and does not require reasoning and computation as complicated as in a one-shot negotiation; and (2) it provides more flexibility for the negotiating parties to efficiently convey information than an ultimatum negotiation protocol, in which one party proposes and the other party can only respond by accepting or rejecting the offers [1, 12]. The negotiation strategy based on an alternating-offers protocol specifies the decisions for both proposal generation and response to an offer.

In a negotiation following an alternating-offers protocol, the negotiators propose and respond alternatively, until one accepts an offer or quits the negotiation¹, or the negotiation deadline T is reached. The history H^t of a negotiation at time t , $t \geq 1$, is a sequence of the negotiators' actions before t , i.e., $H^t = A_i^m_{m < t}$, where A_i^m is the action of negotiator i at time m . Therefore the history of an alternating-offers negotiation at time t is a sequence of proposals, i.e., $H_t = \{x_a^1, x_b^2, x_a^3, x_b^4, \dots, x_a^t(x_b^t)\}$, where x_i^m is the proposal submitted by negotiator i at time m . Generally a negotiation *strategy* S_i specifies the action at each step conditional on the negotiation history, and based on the reservation price and prior belief, i.e., $A_i^t = S_i(H_t | R_i, F_i(\cdot))$, $0 \leq t < T$, where $A_i^t \in \{\text{accept, reject and propose } x_i^{t+1}, \text{quit}\}$. To give an optimal negotiation strategy, game theoretic analysis is required to derive the perfect Bayesian equilibrium² [3]. The analysis of the perfect Bayesian equilibrium is not tractable when both parties have incomplete information with an alternating-offers protocol, although there have been conclusions on the optimal strategy in bilateral negotiations with two-sided incomplete information in a direct revelation mechanism when the prior beliefs of both parties follow a uniform distribution [11][3, Chapter 7]. We adopt two effective negotiation strategies that have been developed in the AI field and proved successful. These two strategies are the *time-dependent strategy* [2] and *Bayesian-learning strategy* [17]. These approaches do not explicitly model the strategic interactions between the ne-

¹A negotiator could quit the negotiation because an agreement is reached in another negotiation thread

²The strategies of players constitute a *perfect Bayesian equilibrium* if given the strategies of the other players, a player cannot obtain strictly better profit *on expectation* in each subgame by deviating to another sequential strategy.

gotiators. Instead they focus on some issues that are important to the decision and information, and provide flexible heuristic decision functions.

2.1.1 Time-dependent approach

The time-dependent approach focuses on the impact of time on negotiations. A negotiator usually has a hard time deadline before which the negotiation has to end. In the Navy detailing system a Sailor has a certain detailing time window which is usually three months. After that, if the Sailor has not located a job he has to be allocated by a detailer. Similarly a Command has to find a Sailor to fill in a job before certain time. With less remaining negotiation time, a negotiator is more pressed to reach an agreement and to concede. But a negotiator cannot wait until the last moment to concede because the time is also valuable. The time spent on negotiation should be reasonable with respect to the value of the agreement that is reached. A negotiator who persistently holds to the price risks losing the negotiation opponent because the opponent may find another partner during the process. There can be other factors that also impact the negotiation strategy such as the available resources. Actually the remaining time can be viewed as one type of resource. The same approach can be used to model the impact of other resources, if there are any, and the decision can add to the decision based on the time-dependent strategy.

In the time-dependent approach time is the predominant factor used to decide which proposal to offer or accept next. For the buyer in a negotiation thread d_i , the proposal to offer or accept is within the interval $[min_i, max_i]$, where max_i is the reservation price of the buyer in thread d_i , and min_i is the lower bound of a valid offer (we can reasonably assume $min_i=0$). For a seller min_i will be the reservation price and max_i is the upper bound of a valid offer. Initially a negotiator offers the most favorable value for herself. If the proposal is not accepted, a negotiator concedes with time proceeding and moves toward the other end of the interval. The pace of concession depends on the negotiation strategy and is characterized by a function $\alpha_i(t)$ of time. The value x_b^t to be offered by a buyer and the value x_s^t to be offered by a seller at time t , $t \in [0, T - 1]$, are as follows:

$$x_b^t = min_a + \alpha_a(t)(max_a - min_a) \quad (1)$$

$$x_s^t = min_b + (1 - \alpha_b(t))(max_b - min_b). \quad (2)$$

The buyer accepts an offer x_s^t from negotiator b at time t if it is not worse than the offer he would submit next time, i.e., $x_a^{t+1} \geq x_b^t$. Similarly the negotiator b accepts an offer x_a^t at time t if $x_b^{t+1} \leq x_a^t$.

The time-dependent function can be defined by a family of polynomial functions³:

$$\alpha_i(t) = \left(\frac{t}{T}\right)^{\frac{1}{\beta}}. \quad (3)$$

The constant $\beta > 0$ determines the concession pace along with time, or the convexity degree of the curve of proposals (see Figure 1). By varying β a wide range of negotiation strategies can be characterized. Two sets of β can be identified to characterize two classes of strategies: Boulware with $\beta < 1$ and Conceder with $\beta > 1$. With a Boulware strategy a negotiator tends to maintain the offered value until the time is almost exhausted, then she concedes to the reservation price quickly. With a Conceder strategy a negotiator goes to the reservation price rapidly and early. Figure 1 shows the change of offers with time in the two strategy classes. Which strategy to use depends on how much a negotiator values the time and the expectation of the opponent's strategy. An impatient negotiator wants to reach a deal earlier and is more likely to follow the Conceder strategy. If a negotiator expects the opponent to be a conceder, she will tend to apply a Boulware strategy.

The time-dependent strategy is intuitive and simple, and has been proved useful in real applications [2]. The shape of the curve of concession, or the parameter β , is what differentiates the strategies of a negotiator. The disadvantage of the approach is that the real-time information in the negotiation is not used. Once β is chosen, the offer curve is pre-determined. But the bilateral negotiation based on an alternating-offers protocol is a sequential interactive process. The information that has been revealed by the opponent in the negotiation can be useful in making subsequent decisions. This consideration is included in the Bayesian-learning strategy.

2.1.2 Bayesian-learning strategy

In the Bayesian learning approach a negotiation agent uses the Bayesian framework to update her prior knowledge and belief about the environment and other agents based on the messages that have been exchanged previously in the negotiation and domain knowledge. Based on her increasingly accurate belief, the negotiator can make better sequential decisions in the negotiation.

³Alternatively we can also use the exponential function family, and define

$$\alpha_i(t) = e^{(1-\frac{t}{T})^{\beta}}.$$

These two families are similar in their functionality except that their sensitivity to the change of time is different with different β . For the same big value of β , the polynomial function concedes faster at the beginning than the exponential one; then they behave similarly. For a small value of β , the exponential function waits longer than the polynomial one before it starts conceding [2].

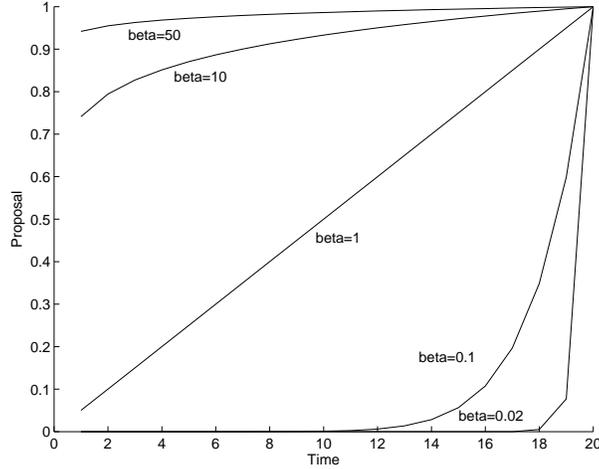


Figure 1: Offer curves with different β

Let the possible type of the opponent be in the collection $\{h_j\}_{j=1}^n$. As we have defined in [6], the *type* of a negotiator is the private information held by the agent that impacts the negotiation outcome. The reservation price is the type of an opponent. The prior belief is the probability that the opponent has a type h_j and is denoted by $P(h_j)$, $j = 1, \dots, n$. The *domain knowledge* attaches a probability $P(e|h_j)$ to every possible action e of the opponent conditional on the type h_j . An opponent proposal action can be viewed as a signal of the opponent's type. Given the encoded domain knowledge in the form of conditional probabilities and the signal e given by the opponent, a negotiator can use the standard Bayesian updating rule to revise her belief about the opponent's type:

$$P(h_j|e) = \frac{P(h_j)P(e|h_j)}{\sum_{k=1}^n P(e|h_k)P(h_k)}. \quad (4)$$

Given her revised belief the negotiator can apply its *decision rule* to make a proposal or respond to an offer. The decision rule can be a simple strategy, for example, (for the buyer) to propose a price which is 10% below the estimated reservation price of the seller. Or it can be a solution to an optimization problem which provides decision heuristics. For example, to make a proposal that maximizes her expected utility assuming the negotiation ends next period. Then the proposal of the buyer at period t is the solution of $\operatorname{argmax}_x F_t(x)(v - x)$, where v is the value of the seller, $F_t(x)$ is the probability that the seller's reservation price is less than x based on the revised belief at period t . The decision of the buyer trades off between the probability of the proposal being rejected and the profit if the proposal is accepted. The solution suggests that

the next proposal x_t satisfies

$$v - x_t = -\frac{F(x_t)}{F'(x_t)}. \quad (5)$$

The domain knowledge can be very specific and confirmative, for example, “in our business a seller usually offers a price 17% above the reservation price”. It can also be simple and “natural” (and cannot be called “domain” knowledge anymore), for example, a seller will not offer a price that is lower than her reservation price. While specific domain knowledge allows efficient update of the prior belief, domain knowledge is hard to identify and acquire. It also requires discretization of the type space to apply the Bayesian framework, as is shown in Equation 4. The “natural” knowledge does not help in the modelling and updating of the prior belief as much as specific domain knowledge, but it is easy to acquire and may not need a discrete type space. For example, let $F_t(x)$ be the current belief at period t on the probability that the reservation price of a seller is less than x , $x \in [\underline{x}, \bar{x}]$. If the seller next proposes a price $z \in [\underline{x}, \bar{x}]$, then the belief can be revised to $F_{t+1}(\cdot)$ based on the knowledge that a seller’s reservation price is always less than her proposal. Then the new belief is:

$$F_{t+1}(x) = \begin{cases} \frac{F_t(x)}{F_t(z)} & \text{for } x \in [\bar{x}, z] \\ F_{t+1}(z) & \text{for } x \in (z, \bar{x}]. \end{cases}$$

Both the belief update method (6) and the decision function (5) have been applied in previous work in negotiations [8, 15], and in other problems such as bidding in double auctions [16, 5]. They can be used to build a basic Bayesian-learning negotiation strategy if no other domain knowledge or more efficient decision rule is available.

2.2 Synchronized multi-threaded negotiations

In a synchronized multi-threaded negotiation process a negotiator participates in multiple bilateral negotiation threads with different, simultaneous negotiation opponents. The negotiator can reach an agreement in at most one of these threads, and is aware of all the threads at the beginning of the process. From one thread’s perspective the other threads are outside options. The reservation utility that the negotiator should set in one thread is equal to the expected utility from all other threads. The other threads form a synchronized multi-threaded negotiation with one less thread than the original process.

As we have explained in [6], with a multi-threaded negotiation it is reasonable to assume that if any agreement is reached, the agreement is signed with the most *competitive* opponents among all opponents of the threads. For a buyer the seller s_i in thread d_i

is more *competitive* than the sellers in other threads if s_i can give more utility to the buyer, i.e., $y_i = v_i - r_i$ is greater than y_j , $d_j \in D \setminus d_i$, where r_i is the reservation price the seller in thread d_i , and $D = \{d_1, \dots, d_N\}$ is the collection of threads. The amount y_i is the *maximum utility* that the buyer can achieve from the negotiation thread d_i . Because the buyer does not know the reservation price of a seller, he does not know the maximum utility in a thread either. Based on the prior belief $F(\cdot)$ on the reservation price of a seller, the negotiator can derive the probability distribution of the maximum utility. From the probability distribution of the maximum utility in every thread, the probability distribution of the highest and second highest maximum utility can be calculated. Let $G_i(y)$ denote the probability of the maximum utility in thread d_i being less than y . Let $G^1(y)$ and $G^2(y)$ be the probability distribution of the highest and second highest maximum utilities. These probabilities can be calculated by the following formulas:

$$G_i(y) = Pr(v_i - r_i \leq y) = Pr(r_i \geq v_i - y) = F(v_i - y)$$

$$G^1(y) = \prod_{d_i \in D} G_i(y)$$

$$G^2(y) = G^1(y) + \sum_{i=1}^N (1 - G_i(y)) \prod_{d_j \in D \setminus d_i} G_j(y).$$

The corresponding probability density functions, or the derivatives of these (cumulative) probability distribution functions, are as follows:

$$g_d(y) = -f(v_d - y)$$

$$g^1(y) = \sum_{d_i \in D} g_i(y) \prod_{d_j \in D \setminus d_i} G_j(y)$$

$$g^2(y) = g^1(y) - \sum_{i=1}^N g_i(y) \prod_{d_j \in D \setminus d_i} G_j(y) + \sum_{i=1}^N (1 - G_i(y)) \left[\sum_{d_j \in D \setminus d_i} g_j(y) \prod_{d_m \in D \setminus \{d_i, d_j\}} G_m(y) \right]$$

We provide four heuristic approaches to estimate the expected utility $OU(D)$ from a multi-threaded negotiation composed by the threads D [6]:

- Conservative estimation: The utility of the buyer is equal to the expected second highest maximum utility. This approach ignores the further concession of the winning seller in the continued single-threaded bargaining process after she outbids the other opponents. The expected utility is calculated by

$$OU = \int_0^{\bar{y}} y g^2(y) dy$$

where \bar{y} is the upper bound of the possible utility that the negotiator can achieve. If the lower bound of an acceptable price for a seller is \underline{c} , and the upper bound of a buyer's valuation is \bar{v} , then $\bar{y} = \bar{v} - \underline{c}$.

- Medium estimation: Assume the continued single-threaded bargaining ends at the middle point between the buyer's and the winning seller's reservation price, if the buyer's reservation price is higher than the winning seller's⁴. Then the expected utility is the average of the expected highest and second highest maximum utility.

$$OU = (\int_0^{\bar{y}} yg^2(y)dy + \int_0^{\bar{y}} yg^1(y)dy)/2$$

In this estimation we do not consider the probability that the negotiation may fail even if an agreement is actually desirable for both parties. This is because in a negotiation model with incomplete information, negotiators are not willing to reveal their reservation prices but expect the concessions of the other. This inefficiency is considered in the approach of uniform approximation.

- Uniform approximation: Previous research has established an optimal bargaining result between a buyer and a seller based on game theoretic analysis when both parties' reservation prices follow uniform distributions [11, 3]. Based on this result, an agreement occurs if and only if the buyer's valuation exceeds the seller's cost by at least 1/4, if both parties' reservation prices distribute uniformly on [0, 1]. In other words, an agreement cannot be reached if the buyer's valuation is less than the seller's cost plus 1/4 of the maximum difference between the buyer's valuation and the seller's cost. We can approximate the probability distributions of negotiators' types by uniform distributions and apply this result to calculate the probability of reaching an agreement. In the heuristic we assume an agreement cannot be reached in the continued single-threaded negotiation between the buyer and the winning seller if the maximum utility of the winning seller is less than a quarter of the highest possible utility \bar{y} . In this case the buyer achieves the second highest maximum utility, which is the reservation utility of the buyer in the continued single-threaded negotiation. If an agreement is reached in the single-threaded negotiation, it is reasonable to assume that it is at the middle point between both parties' reservation prices. Therefore in this case the buyer achieves the medium of the highest and the second highest maximum utility.

$$OU = \frac{\int_0^{\bar{y}} yg^2(y)dy + \int_0^{\bar{y}} yg^1(y)dy}{2} \int_{\bar{y}/4}^{\bar{y}} g^1(y)dy + \int_0^{\bar{y}} yg^2(y)dy(1 - \int_{\bar{y}/4}^{\bar{y}} g^1(y)dy).$$

⁴If the buyer's reservation price is lower than the seller's, there is no "zone of agreement" and the negotiation will fail.

- Learning: Learn the probability of reaching an agreement and the distribution of agreements based on the previous negotiations [14]. The result of learning is represented by $x(R_b, R_s)$, the expected agreement on the price from the negotiation when the buyer's and seller's reservation prices are R_b and R_s respectively. Given the probability distribution of the opponent's reservation price, a negotiator can calculate the expected utility of the negotiation based on the result of learning. If the seller s_i in the thread d_i is the winning seller, then the probability distribution of her reservation price is $F(c) \prod_{d_j \in D \setminus d_i} (1 - F(v_j - v_i + c))$, where the product is the probability that no other thread d_j has the maximum utility $v_j - c_j$ greater than the maximum utility $v_i - c_i$ in thread d_i . Then the expected utility from a multi-threaded negotiation can be calculated by

$$OU = \sum_{d_i \in D} \int_{\underline{c}}^{\bar{c}} (v_d - x(v_i, c)) \prod_{d_j \in D \setminus d_i} (1 - F(v_j - v_i + c)) dF(c)$$

If negotiators use the time-dependent strategy and the parameter β is chosen randomly with the mean equal to 1, then we expect negotiators to concede constantly on aggregation. Then the result of learning is expected to be close to the result of negotiation when $\beta = 1$ for both negotiators. Let the reservation prices of the buyer and the seller be v and c respectively. Then

$$x(v, c) = \begin{cases} \frac{v}{1+v-c} & \text{if } v \geq c \\ \emptyset & \text{otherwise} \end{cases} \quad (6)$$

assuming the upper bound of an offer is 1 and the lower bound is 0. This formula can be derived from Figure 2 that shows the offer curves of the negotiators with $\beta = 1$ for both parties:

2.3 Dynamic multi-threaded negotiations

In the Navy detailing process the application period for a position, or search period for filling a job, lasts for some months. During that period potential partners are discovered sequentially and new negotiations are launched dynamically. For an ongoing negotiation thread the outside options not only include the other simultaneous negotiation threads, but also the threads that may be launched in the future. Considering the outside options in the future, a negotiator must decide how much to offer in the current negotiation, and when to stop searching for future opportunities and accept an offer from the current negotiation. If a negotiator knows the number of outside options that will come, and the value of the opponent in each outside option, then the

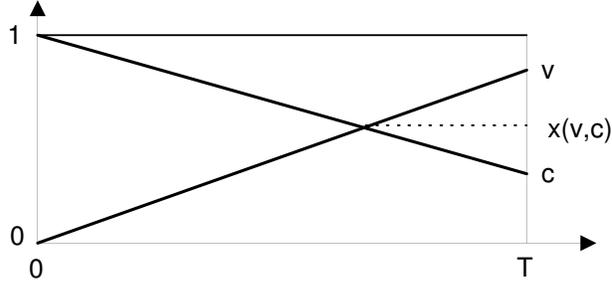


Figure 2: Offer curves with $\beta = 1$

negotiator can apply the synchronized multi-threaded negotiation model to calculate the appropriate reservation price in each thread. But in the Navy detailing process, neither a Command nor a Sailor is sure about the arrival of, and the opponents' values in, future outside options. The reservation utility of a thread is the expected utility of a multi-threaded negotiation - including other simultaneous threads and threads launched in the future - with a stochastic thread number and uncertain opponents.

Following the usual way of modelling uncertain arrival, we assume the arrival of outside options follows a Poisson process [9, 10, 13]. In each period t , $t = 0, \dots, T - 1$, there is probability p that the buyer finds a matching alternative and launches a negotiation thread. The granule of each period is small enough so that the probability that there are more than one arrival in one period is zero. Again the value v of a seller follows the probability distribution $\Phi(y) = Pr(v \leq y)$, where $\Phi(y)$ is the probability that a seller's value is no greater than y . The reservation price R_s of a seller follows the prior belief $F(x) = Pr(R_s \leq x)$. A Command can evaluate a Sailor by checking the Sailor's background. A Sailor also knows how much he prefers a job by acquiring the job information about location, responsibility, etc. But how much a Command values a Sailor or a Sailor values a Command is unknown to the Sailor or the Command respectively. Therefore a negotiator knows the value of an opponent when the opponent is identified, but not the reservation price of the opponent.

The *state* s_t of the system is defined as the number of threads n_t and the value of each opponent seller v_d , $s_t = \{n_t, \{v_d\}_{d=1}^{n_t}\}$. The evolution of the system follows the rule

$$s_{t+1} = \begin{cases} \{n_t + 1, \{v_d\}_{d=1}^{n_t} \cup v\} & \text{if an opponent with value } v \text{ arrives at period } t \\ s_t & \text{if no arrival at period } t \end{cases}$$

Let $OU_t(s_t)$ be the utility that the negotiator expects from the dynamic multi-threaded negotiation when she sees the system state s_t at period t . Following Section 2.2 we can calculate $OU(\{n, \{v_d\}_{d=1}^n\})$, the expected utility from a synchronized multi-threaded negotiation with n threads and the opponent in thread d valued v_d , $d = 1, \dots, n$. The transition of the expected utility follows the rule

$$\begin{aligned} OU_t(s_t) &= (1 - p)OU_{t+1}(s_t) + pE_v[OU_{t+1}(\{n_t + 1, \{v_d\}_{d=1}^{n_t} \cup v\})], \\ OU_{T-1}(s_{T-1}) &= OU(s_{T-1}). \end{aligned} \quad (7)$$

If the probability of arrival at each period is p , then the number of arrivals $\eta(m, p)$ during an interval with length τ follows a Poisson distribution, $P_{p,\tau}(n) = Pr(\eta(\tau, p) = n) = e^{-p\tau} \frac{(p\tau)^n}{n!}$. Equivalently we can write the transition of the expected utility as

$$OU_t(s_t) = E_\eta[E_{\{v_d\}_{d=n_t+1}^{n_t+\eta}} [OU(\{n_t + t, \{v_d\}_{d=1}^{n_t} \cup \{v_d\}_{d=n_t+1}^{n_t+\eta}\})]] \quad (8)$$

where η following a Poisson distribution $P_{p,T-t}(\cdot)$, and v_d independently follows the identical distribution $\Phi(\cdot)$, $d = n_t + 1, \dots, n_t + \eta$.

To set the reservation price of a thread, the negotiator only needs to calculate the expected utility of the multi-threaded negotiation which does not include that thread, based on the period and real-time state. Because the state of a dynamic multi-threaded negotiation changes from period to period, the reservation price of a thread may also change with time.

The expected utility of a dynamic multi-threaded negotiation process at each period with each state can be calculated backward from the last period following Equation 7 or forward following Equation 8. The computation will be very expensive to calculate the expectation of the expected utility on the opponents' values. If there are at most N threads and for each opponent there are M possible values, then the number of possible states will be N^M . The computation is intractable with large M . To simplify the computation we can approximate the result by having the opponent value instances replaced by the expected value \bar{v} , i.e.,

$$OU_t(s_t) = (1 - p)OU_{t+1}(s_t) + pOU_{t+1}(\{n_t + 1, \{v_d\}_{d=1}^{n_t} \cup \bar{v}\}) \quad (9)$$

which is equivalent to

$$OU_t(s_t) = E_\eta[OU(\{n_t + t, \{v_d\}_{d=1}^{n_t} \cup \{\bar{v}\}_{d=n_t+1}^{n_t+\eta}\})] \quad (10)$$

The compromise due to this simplification is not significant if the expected utility of a synchronized thread is or can be approximated by a linear function of the opponents' values.

3 Experiments

In Section 2 we have presented two models of the outside options, the synchronized and dynamic multi-threaded negotiation models, and four heuristic approaches, the conservative estimation, the medium estimation, the uniform approximation and learning approach, to estimate the expected utility in a multi-threaded negotiation. By combining different outside option models and estimation approaches, we can have eight decision models for bilateral negotiations in the Navy detailing process. In this section we provide experiments to illustrate the different models in the solution framework and the performance results based on the different models. We used the time-dependent strategy as the strategy in a single-threaded negotiation because it is simple to compute. In the solution framework that we have presented, the reservation utility is an important system variable that decides the reservation price, which impacts the offer curve based on a specific negotiation strategy. In Section 3.1 we show how the reservation utility of a negotiation thread evolves with time and the change of outside options in the synchronized and dynamic multi-threaded negotiation models. We then show the impact of outside options on the negotiation strategy by showing the offer curves adjusted by the reservation prices, compared with the original basic offer curve without considering outside options. In Section 3.2 we compare the average utility of a negotiator when she (1) does not consider outside options, (2) when she only considers concurrent outside options, i.e., the synchronized multi-threaded negotiation model, and (3) when she considers both concurrent outside options and future arrivals, i.e., the dynamic multi-threaded negotiation model. The performance results based on different utility estimation approaches are also compared and discussed.

3.1 Reservation utilities and offer curves

We illustrate the impact of outside options on the negotiation strategy by a specific example. In this example the negotiator is a buyer. The negotiation deadline $T = 20$. The buyer believes the reservation price of a seller follows a uniform distribution on the interval $[0, 1]$. The value of (the item of) a seller to the buyer is also uniformly distributed on $[0, 1]$. In each period with probability $p > 0$ a new seller arrives and a negotiation thread is created. The shape of the offer curve defined by Equation 3 and 1 (Section 2.1) is determined by the parameter β . We ran multiple experiments with $p = 0.05, 0.10, 0.15, 0.20, 0.25$ and random β . The resulting curves followed the same pattern for all instances. We show the resulting curves based on one instance with $p = 0.2$ and $\beta = 1.262727$. The arrivals of outside options in the instance are illustrated in Figure 3. The figure on the left shows the time of arrivals and the value of each arrival, the figure on the right shows the number of threads in each period.

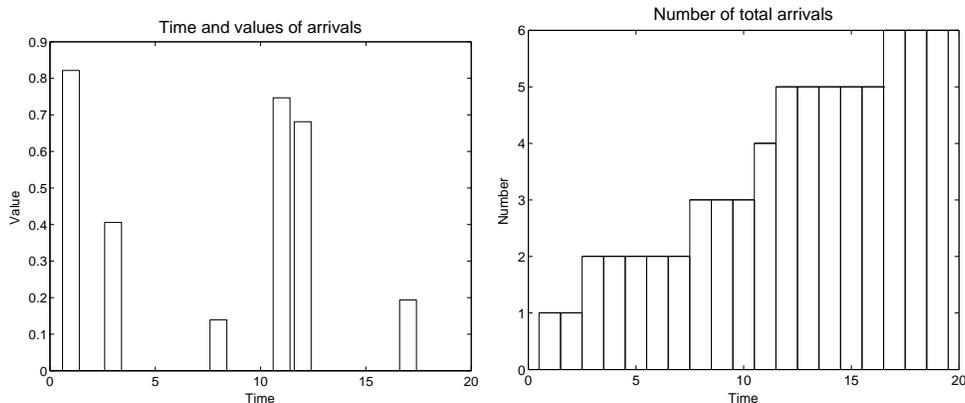


Figure 3: Arrivals and number of threads

To illustrate the evolution of the reservation utility of a thread, we collect the reservation utility of the first thread along time calculated with different estimation approaches and models of the outside options. Figure 4 shows the reservation utilities grouped by the estimation approach and Figure 5 compares the reservation utilities calculated based on different estimation approaches but on the same outside option model. The expected utility of a dynamic multi-threaded negotiation process was calculated with the approximation formula, Equation 10 (Section 2.3).

We can see from Figure 4 that the reservation utility based on the synchronized model without expecting arrivals in the future is less than the reservation utility calculated based on the dynamic model where the probable future arrivals are also taken into consideration⁵. This is as expected because in the dynamic model a negotiator does not only see the current existing outside options, but also foresees the outside options that may come in the future. Although the number, the arrival times and opponents'

⁵There is a region in which the reservation utility based on the synchronized model is slightly higher than the one based on the dynamic model. This is because of the learning model that is used. Consider the situation where there is only one thread d_1 . The expected utility based on the learning model is $v_1 \frac{v_1 - c_1}{1 + v_1 - c_1}$. Now assume a new negotiation opponent comes, and the value of the opponent is v_2 and the reservation price is c_2 . The estimation model suggests that the expected utility from the two-threaded negotiation is $v_2 \frac{v_2 - c_2}{1 + v_2 - c_2} Pr(v_2 - c_2 > v_1 - c_1) + v_1 \frac{v_1 - c_1}{1 + v_1 - c_1} Pr(v_1 - c_1 \geq v_2 - c_2)$. When v_1 is much greater than v_2 , $v_2 \frac{v_2 - c_2}{1 + v_2 - c_2}$ may be less than $v_1 \frac{v_1 - c_1}{1 + v_1 - c_1}$ even if $v_2 - c_2$ is greater than $v_1 - c_1$. Therefore more threads do not necessarily mean higher expected utility. Arrivals with very low values may actually reduce the expected utility calculated by the learning approach. But generally we can say that the expected utility increases with the number of threads and hence the expected utility based on the dynamic model is higher than the one based on the synchronized model.

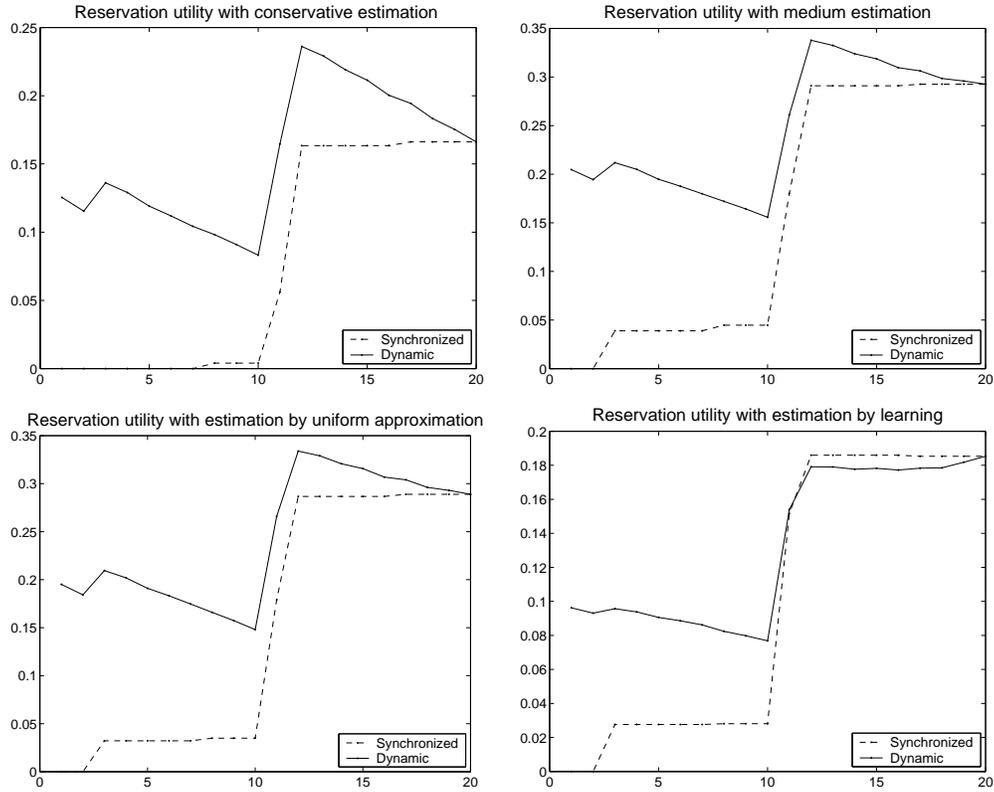


Figure 4: Reservation utilities grouped by estimation approaches

values of the future outside options are uncertain, they are still valuable and provide opportunities of reaching an agreement outside the current negotiation threads. Therefore the expectation on the possible future outside options raises the lowest utility that a negotiator can accept in a current negotiation thread.

Figure 5 shows that the approach of medium estimation suggests a higher reservation utility than the other approaches and gives the most optimistic estimation on the utility from outside options. The medium estimation gives a more optimistic estimation than the conservative estimation because in the latter the concession of the winning opponent in the continued single-threaded negotiation is ignored. The expected utility based on the medium estimation is higher than on the uniform approximation estimation because the inefficiency of a negotiation with two-sided incomplete information is considered in the latter but not in the former. The medium estimation also suggests a higher expected utility than the learning approach because in the latter the negotiation

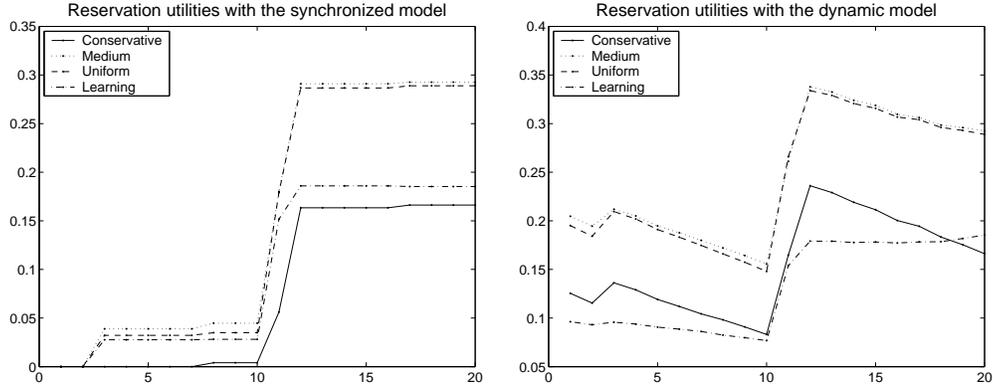


Figure 5: Reservation utilities grouped by outside option models

outcome in the continued single-threaded negotiation is not compared with the second highest maximum utility, while in the former it is guaranteed that the negotiation outcome in the continued single-threaded negotiation is not worse than the second highest maximum utility.

No matter which estimation approach is used, Figure 5 shows that the reservation utility based on the synchronized model (Section 2.2) monotonically increases with time because the number of threads increases with time. But it is interesting to note that the reservation utility based on the dynamic model (Section 2.3) is not a monotonic function of time. This is because there are two forces that drive the change of the reservation utility: time and concurrent threads. When the negotiator approaches the deadline, the possibility to have new arrivals decreases and it drives the reservation utility down. On the other hand the reservation utility would increase with the arrival of a new negotiation opponent, especially when the value of the new opponent is high. From Figure 5 we can see that the reservation utility has different sensitivity to the change of time and new arrivals based on different estimation approaches. The reservation utility based on the learning approach does not change as much as the ones based on other approaches as time or outside options vary. No matter which estimation approach or outside option model is used, the resulting reservation utility with consideration of future outside options is higher than without considering the future outside options.

Based on the reservation utility OU_d of thread d we can calculate the reservation price of the buyer, $R_d = v_d - OU_d$, where v_d is the value of the seller in thread d . We compare the offer curves based on different estimation approaches and outside option

models in Figure 6. The model noted by “Single” is the model without considering outside options. In that model the reservation price R_d is equal to the opponent’s value v_d and the offer x_t in period t is calculated by $x_t = (t/T)^{1/\beta}v_d$ following Equation 1 (We assume the lower bound min_b of a valid offer is zero.). Without considering outside options, the offer increases with time as the buyer constantly concedes (with changing pace). But with a synchronized or dynamic model the buyer may proceed, i.e., decrease the offered price from the previous one, when a valuable new opponent arrives. The pace of concession is also different with different outside options models. When a valuable seller arrives, the buyer may proceed by asking for a lower price than the previous offer because the buyer gets more optimistic about the expected utility that she can get from the outside options. In the dynamic model the speed of concession at some time may be higher than without considering outside options in the single model, because of the impact of both the increasing time pressure on reaching an agreement and the decreasing hope on the availability of future outside options. The offers without considering outside options are higher than the offers with considering only concurrent negotiation threads, which are again higher than the offers with additional considerations of outside options that may come in the future. This is consistent with the observation that the reservation utility based on the synchronized model is less than the one based on the dynamic model.

3.2 Performance results

In this section we examine and compare the average utilities that a buyer obtains with three different outside option models and four different estimation approaches. The three outside option models include: (1) the “Single” model in which no outside options are considered, (2) the “Synchronized” model in which only concurrently existing negotiation threads are considered as outside options, and (3) the “Dynamic” model in which the outside options also include the possible uncertain future arrivals. The four estimation approaches include the conservative estimation, the medium estimation, the uniform approximation and the learning approach. In the experiments the negotiation deadline $T = 20$. The buyer believes the reservation price of a seller follows a uniform distribution on the interval $[0, 1]$. The value of a seller’s item is also uniformly distributed on $[0, 1]$. The probability that a new seller arrives in a period is p , and p takes the values $\{0.05, 0.10, 0.15, 0.20, 0.25\}$. The parameter β in the time-dependent strategy of a negotiator is chosen randomly so that with even probability a negotiator in a thread is a conceiver ($\beta > 1$) or a brawler ($\beta < 1$). If a negotiator is a conceiver, β^{-1} follows a uniform distribution on $[0, 1]$. If a negotiator is a brawler, β is a random variable with a uniform distribution on $[0, 1]$. For each arrival probability, we repeat the experiment 50 times and the average utility of the buyer is calculated.

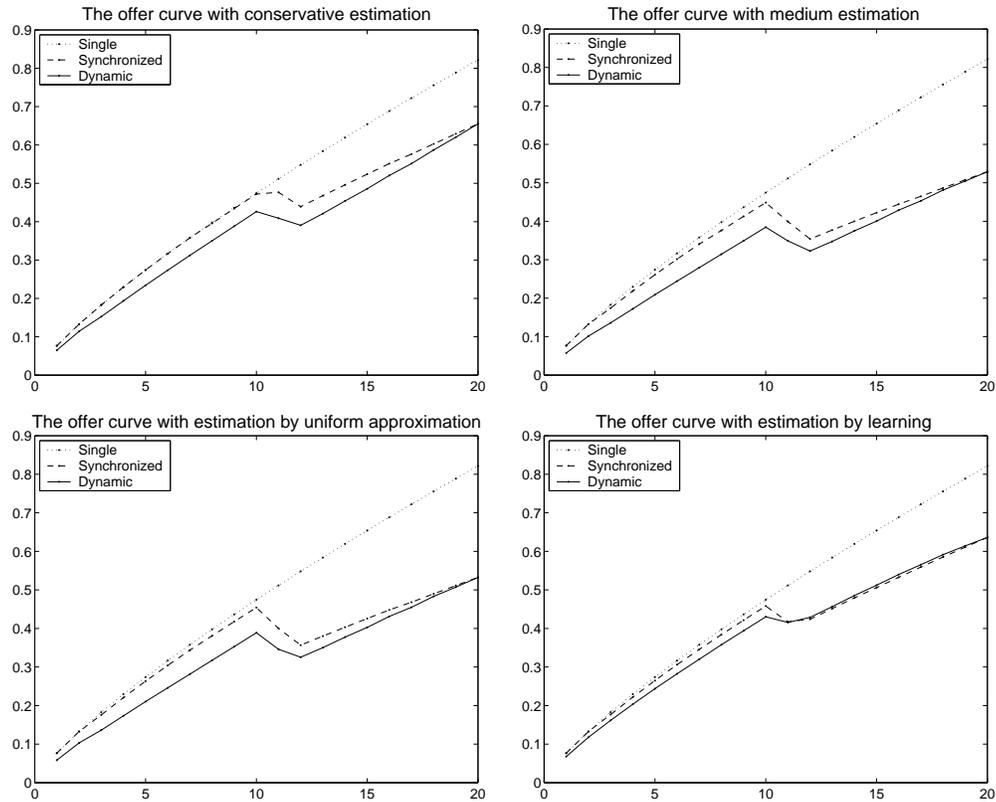


Figure 6: The offer curves

Figure 7 is composed of four subplots. Each subplot shows the average utility as a function of the arrival probability based on one estimation approach, and with different outside option models. The figure implies that for all estimation approaches and outside option models, the average utility increases with the arrival probability. This is intuitive and should be true for a reasonable negotiation strategy. A higher arrival probability implies more options on expectation and should result in better outcome for the negotiator. Figure 7 also shows that the average utility based on the dynamic model is higher than the one based on the synchronized model, which again brings higher average utility than the single-threaded model in which no outside option is considered. This verifies the effectiveness of the outside option models that we have proposed.

We can also group the average utilities by the outside option model and compare the performance of the estimation approaches. The information is shown in Figure 8. The

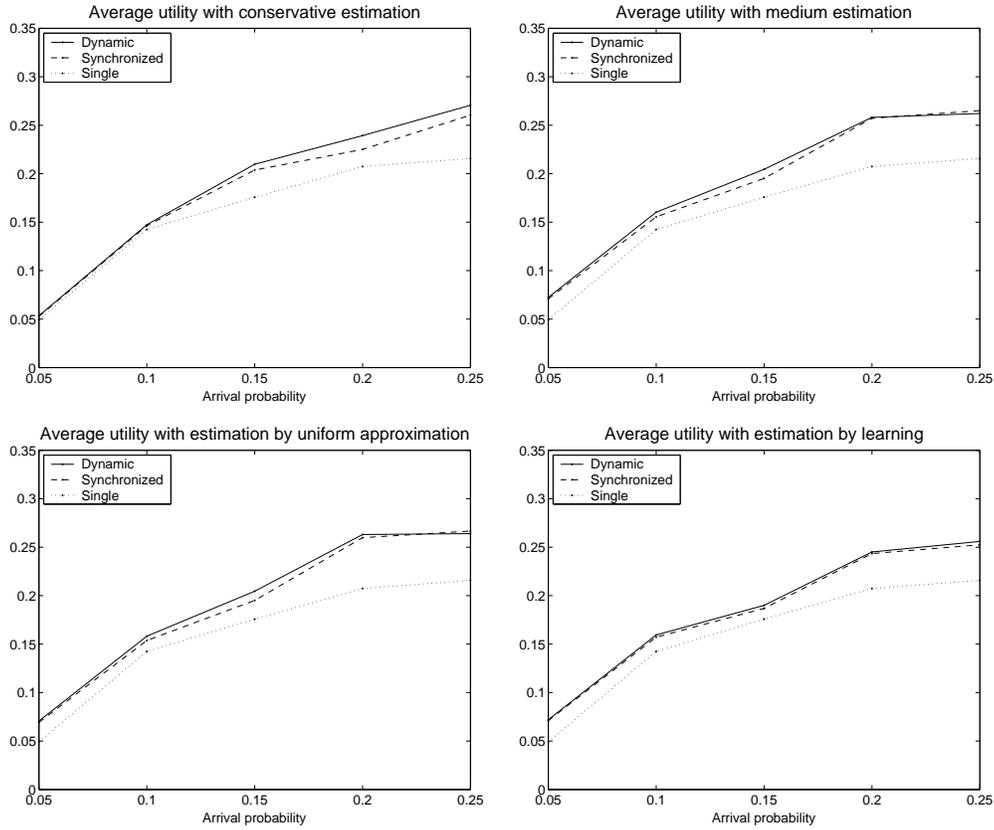


Figure 7: The average utilities grouped by estimation approaches

figure shows that there is no estimation approach that dominates the others. This is because the performance of an approach depends on the negotiators' offer curves. If both negotiators tend to concede quickly (β is very large), an optimistic estimation approach such as the medium approach may be better. On the other hand if both negotiators tend to hold on their positions (β is very small), the conservative estimation approach may be better.

4 Discussion and conclusion

In this report we provide an integrative solution for the negotiation decision problem in the Navy detailing system when negotiators face uncertain and dynamic outside

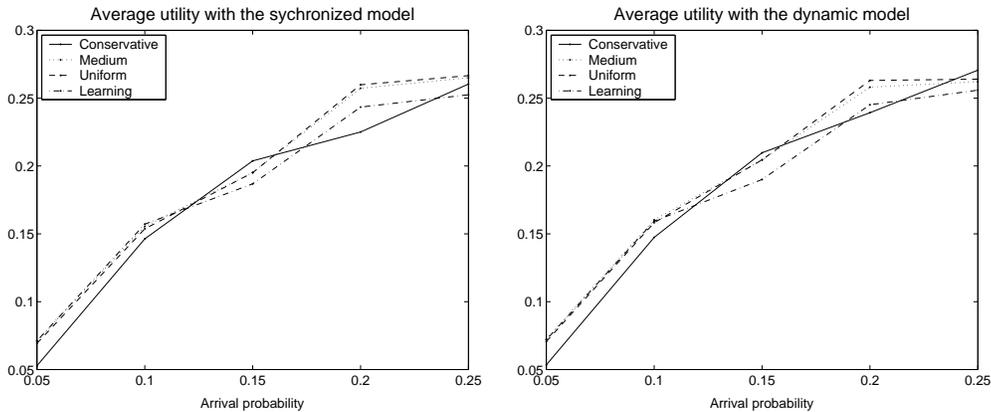


Figure 8: The average utility grouped by outside option models

options. The outside options influence the negotiation strategies via the impact on the reservation prices. The solution is composed of three modular models: single-threaded negotiations, synchronized multi-threaded negotiations, and dynamic multi-threaded negotiations. The single-threaded negotiation model provides the negotiation strategy given the reservation price. The other two models calculate the reservation price based on the model of the outside options. The model of synchronized multi-threaded negotiations considers the presence of concurrently available outside options and provides an approach to estimate the outcome when the threads are known. The model of dynamic multi-threaded negotiations expands the last model by considering the uncertain outside options that may come dynamically in the future. The specific solution for each module is presented, and experimental analysis is provided. The results show that the utility of a negotiator improves significantly when she considers outside options than when she does not consider them, and when she considers the dynamic arrival of outside options than when she only considers the concurrently existing negotiation threads.

We would like to make the following remarks to avoid possible confusions in understanding the model:

- We take an artificial intelligence approach instead of a game theoretic or economics approach in this study because the complexity of the situation does not allow rigorous mathematical analysis, which is usually preferred by economists. The approaches in the AI field are different from the models in economics in that AI approaches aim to provide an effective heuristic solution to general, realistic

and complicated situations, since in complicated models there is lack of rigorous mathematical analysis.

- We are not using an auction mechanism⁶. In a multi-threaded negotiation process, each negotiation thread is an outside option of other threads. But in an auction which includes all candidates as bidders, there is no outside option for the auctioneer.
- The "reservation price" should not be confused with the "intrinsic value" of an item. The reservation price from a buyer's point of view is the highest price that is acceptable in one negotiation. The highest acceptable price is no greater than the intrinsic value. When there are outside options, the highest acceptable price may be less than the intrinsic value.
- Outside options change the reservation price in a negotiation. To give a very simple example, if a buyer knows that she could buy an item from a seller for \$20, she will not buy it from another seller for more than \$20, although she values the item at \$25. In the Navy situation, a buyer is uncertain about the availability and quality of outside options, neither does she know the exact agreement she could reach in other outside negotiations. She could only set an appropriate reservation price by estimating the utility she could achieve from outside options with reasonable heuristic approaches, which we have provided in our deliverable.

In this negotiation solution we have focused on the negotiation strategy when the negotiator faces uncertain outside options. We did not explicitly model the behavior of the negotiation opponents when they also have outside options. The outside options of an opponent are unknown to the negotiator and influence the reservation price of the opponent. Since the reservation price is private information, the outside options of an opponent can be taken into consideration if the prior belief on the opponent's reservation price also includes the probabilistic information on her outside options.

In this report we propose heuristic solutions for the negotiation decision problem in the Navy detailing process. The complexity of interactions in an alternating-offers bilateral negotiation with two-sided asymmetric information does not allow the mathematical analysis of optimal strategies with general settings such as continuous type space and general probability distribution of the prior belief, even without outside options. In this work we pursue the practical effectiveness of the solution. We have proposed and applied negotiation strategies that have been developed by us and others in the AI field, and we have provided reasonable heuristic approaches to set appropriate reservation

⁶Even with auctions, the reservation price is not necessarily equal to the true value [Auction Theory, Vijay Krishna, Academic Press, 2002].

prices considering outside options. Existing results from economics on the optimal solutions of simpler models, such as auctions and bilateral negotiations with uniform distributions of the prior beliefs, are used to design reasonable heuristics to solve the complicated problem in our model. Because of the heuristic approach that we take in this work, extensive simulations are needed to provide rigorous evaluation on the performance of different models and approaches in different environments. We do not claim that the heuristics we provide in this report are complete. Rather they reflect solutions that have been proven useful or plausible. Other negotiation strategies and approaches to estimate the utility from a multi-threaded negotiation can be plugged in the solution framework, depending on the assumptions and requirements of the underlying application. These different models can construct a library of decision functions to support the decision of negotiation agents in different environments.

Bilateral negotiations are a useful mechanism to realize distributed matching between Sailors and Commands that do not get matched through the mass-matching market. How these two mechanisms interact with each other depends on Navy policy. If a Command or a Sailor has to accept the matching result from the mass-matching market, then the bilateral negotiations can be regarded as outside options of the mass-matching market. How much to bid in the mass-matching market is impacted by the reservation utility that a Command or a Sailor expects to obtain via bilateral negotiations. Otherwise if a Command or a Sailor can reject the results from the mass-matching market, then the bidding decision is not entangled with the bilateral negotiation process. But the decision to accept or reject the outcome of the mass-matching mechanism still depends on the expected utility from the bilateral negotiations.

Acknowledgements

This work was supported in part by the Navy Personnel Research, Studies and Technology (NPRST) under Navy contract number N6610-98-D-9501 and by a subcontract with Resource Consultants, Inc (RCI), subcontract number 03-177-7048/086.

References

- [1] Lawrence Ausubel, Peter Cramton, and Raymond Deneckere. *Bargaining with incomplete information*, volume 3, chapter 50. Elsevier Science, 2002. Robert J. Aumann and Sergiu Hart, eds.

- [2] 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.
- [3] Drew Fudenberg and Jean Tirole. *Game Theory*. MIT Press, 1991.
- [4] Joseph A. Giampapa, Cuihong Li, Wei Yang, and Katia Sycara. Recommendations and descriptions of cognitive agents for the navy detailing task. Technical Report CMU-RI-TR-03-31, The Robotics Institute, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213-3890 (U.S.A.), August 2003. Version 1.1.
- [5] S. Gjerstad and J. Dickhaut. Price formation in double auctions. *Games and Economic Behavior*, 1:1–29, 1998.
- [6] Cuihong Li, Joseph Giampapa, and Katia Sycara. Modelling bilateral negotiations in the Navy detailing process. Technical Report CMU-RI-TR-03-47, Carnegie Mellon University, 2003.
- [7] Cuihong Li, Joseph Giampapa, and Katia Sycara. A review of research literature on bilateral negotiations. Technical Report CMU-RI-TR-03-41, Carnegie Mellon University, 2003.
- [8] Cuihong Li and Gerald Tesauero. A strategic decision model for multi-attribute bilateral negotiation with alternating offers. In *Proceedings of the 4th ACM Conference on Electronic Commerce*, pages 208–209, June 2003.
- [9] S.A. Lippman and J.J. McCall. The economics of job search: A survey. *Economic Inquiry*, XIV:155–189, 1976.
- [10] S.A. Lippman and J.J. McCall. *The economics of uncertainty: Selected topics and probabilistic methods*, volume I, chapter 6. North-Holland Publishing Company, Amsterdam, in k.j. arrow and m. intriligator, editors edition, 1981.
- [11] Roger B. Myerson and Mark A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29:265–281, 1983.
- [12] Stefan Napel. *Bilateral Bargaining: Theory and Applications*. Lecture notes in economics and mathematical systems. Springer, 2002.
- [13] Sheldon M. Ross. *Introduction to Probability Models*. Academic Press, 2000.
- [14] Katia Sycara. Machine learning for intelligent support of conflict resolution. *Decision Support Systems*, 10:121–136, 1993.

- [15] Gerald Tesauro. Efficient search techniques for multi-attribute bilateral negotiation strategies. In *Proceedings of ISEC-02*, 2002.
- [16] Gerald Tesauro and Jonathan L. Bredin. Strategic sequential bidding in auctions using dynamic programming. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2002.
- [17] Daniel D. Zeng and Katia Sycara. Bayesian learning in negotiation. *International Journal of Human-Computer Studies*, 48:125–141, 1998.