Bilateral Negotiation Decisions with Uncertain Dynamic Outside Options

Cuihong Li, Joseph Giampapa and Katia Sycara

***Abstract*— We present a model for bilateral negotiations that considers the uncertain and dynamic outside options. Outside options affect the negotiation strategies via their impact on the reservation price. 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 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 existing outside options. 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. Experimental analysis is provided to characterize the impact of outside options on the reservation price and thus on the negotiation strategy. The results show that the utility of a negotiator improves significantly if she considers outside options, and the average utility is higher when she both considers the concurrent outside options and foresees future options.**

I. Introduction

Bilateral negotiations are important mechanisms to achieve distributed conflict resolution when it is in the common interest of the parties to cooperate [1], [2]. Automated negotiation strategies have been one of the most fundamental decision models for implementing self-interested and autonomous interacting software agents [3], [4]. A negotiation strategy is a mapping from input information about the *environment* to a sequence of decisions. The environment includes all factors that impact the negotiation outcome, for example, valuations of agreements, the possible valuation held by the "opponent"[1], and the deadline for reaching an agreement. With given inputs (e.g., reservation prices and the deadline), design of the mapping, or the negotiation strategy, is important for ensuring an efficient negotiation outcome. On the other hand, it is usually not straightforward to acquire the input information for a negotiator situated in an uncertain and dynamic environment. Modeling the environment and the impact of the environment is important for enabling an agent to conduct an efficient negotiation, both reactively and proactively, based on the change, and the prediction of the change of the environment.

Usually a negotiator can face more than one candidate to reach an agreement with, although only one agreement with a single candidate is allowed. These candidates become *outside options* with respect to each other for the negotiator. The outside options contribute to the environment of the negotiation with a candidate. A motivating example is the matching market in the Navy detailing system, which allocates sailors on rotation to job vacancies (billets). Bilateral negotiations between sailors and commands are one of the mechanisms proposed to achieve distributed detailing so that personal preferences, such as the location and training opportunities, can be considered, and matching incentives, such as vacation time and payment, can be provided. As a sailor approaches the end of his/her current duty, the detailing system reacts to this by creating a vacancy and putting this vacancy on the matching market. A sailor enters a rotation period and displays on the matching market three months before the end of his/her current tour. This situation occurs thousands of times during a year. A sailor may find multiple jobs that he/she is qualified for and interested in, and similarly a Navy command in charge of a job may see multiple sailors available and qualified for that position. Since the matching market at any time only contains the information that reflects available sailors and vacancies in the next few months, a command and a sailor also expect the possible arrival of more alternatives in the future of the rotation time window. The outside options are an important issue to consider in a negotiation between a command and a sailor for fulfilling a job vacancy.

Existence of outside options is typical in matching markets, and also common in commodity and service markets [5], [2]. Accepting a proposal in one negotiation means refusing all outside options. On the other hand one may leave a negotiation before the deadline (called "opt-out" of a negotiation) without reaching an agreement based on the expectation of reaching a more favorable agreement in outside options. We call the negotiation between a negotiator and one of the opponents a negotiation *thread*. Modeling the outside options and understanding the interaction between outside options and a negotiation thread is an essential aspect to designing an effective negotiation strategy in the environment with outside options. For convenience of presentation, in the rest of the paper we call the two parties in a negotiation a buyer and a seller, and the model is presented from a buyer's perspective. A similar model can be built from a seller's perspective.

Outside options can exist *concurrently* with a negotiation thread, or come *sequentially* in the future. A *concurrently*

Cuihong Li is a PhD candidate at Tepper School of Business in a joint program with Robotics Institute, Carnegie Mellon University. Email: cuihong@andrew.cmu.edu

Joseph Giampapa is a researcher and the project manager of the Intelligent Software Agents lab in Robotics Institute, Carnegie Mellon University. Email: garof+@cs.cmu.edu

Dr. Katia Sycara is a professor and the director of the Intelligent Software Agents lab in Robotics Institute, Carnegie Mellon University. Email: katia@cs.cmu.edu

[1] Here we refer to the other party against whom a party is negotiating as the "opponent".

existing outside option is a negotiation thread that the negotiator is involved in simultaneously with another thread. This happens because a buyer may find multiple sellers that provide the desired item and are available for negotiations at the same time. A *sequentially* available outside option is an exchange opportunity that comes in the future. In a dynamic environment sellers may not come to the market or be observed by the buyer at the same time.

Outside options are *uncertain* in terms of both *availability* and *quality*. The *availability* of outside options is uncertain because the the buyer is not sure when an outside option is available and how many are available. The *quality* of outside options is uncertain because the buyer cannot predict the value of the item provided by a seller who comes in the future, non the outcome of a negotiation thread. The items provided by different sellers are heterogeneous in quality. A buyer does not know the value of an item until she meets the seller and sees the item.

Outside options impact the input to a negotiation decision model as a part of the environment. The existence of outside options changes the utility that the buyer expects from a negotiation thread, and hence the agreement that is acceptable for the buyer in the thread [6]. We claim that *outside options affect the negotiation strategies via their impact on the buyer's reservation price*. The *reservation price* is the worst agreement that a negotiator can accept. For the buyer the reservation price is the highest price she is willing to pay for the negotiated good. The buyer's reservation price depends on the quality/value of the good provided by the seller, and also on the availability of other sellers. The buyer is not willing to accept a price from a seller if that price brings a lower utility than she expects from outside options. In other words, the buyer expects a utility from a thread that is not lower than from outside options. The expected utility from outside options becomes the *reservation utility*, the least utility that is acceptable in a negotiation thread. The reservation utility determines the reservation price of the buyer in a negotiation thread, which again impacts the proposal and response strategies.

The design of an effective negotiation strategy can be divided into two parts: the first part is the design of a negotiation strategy given the reservation price and other inputs, the second part is to calculate the reservation price based on the model of outside options. We call the model in the first part *single-threaded negotiations*.

The model of outside options can be built with two levels of complexity based on the two forms of availability of outside options. On the first level the buyer assumes there are no outside options coming in the future, and makes decisions based on the outside options that concurrently exist with the thread under consideration. Therefore there is no uncertainty about outside options in terms of both the number of threads and item values. When new outside options arrive, the buyer makes corresponding adjustments to her reservation price and negotiation strategy *reactively*. We call this model on negotiations with only concurrently available outside options *synchronized multi-threaded negotiations*. On the second level the buyer also considers the outside options that may come dynamically in the future. Hence in the decision model the number of threads that the negotiator would be involved in is a random variable and changes with time. The item values in the future threads are also uncertain. The buyer acts both *reactively* to the realized outside options, and *proactively* to the possible arrivals based on the prediction information. We call this model with both concurrently and sequentially available outside options *dynamic multi-threaded negotiations*. It builds on the synchronized multi-threaded model but introduces uncertainty to the threads. In both models of synchronized and dynamic multi-threaded negotiations the negotiation strategy in one thread can be derived from the single-threaded negotiation model, but the reservation price is calculated with the corresponding model of outside options. Figure 1 shows the relationship between these three negotiation models.

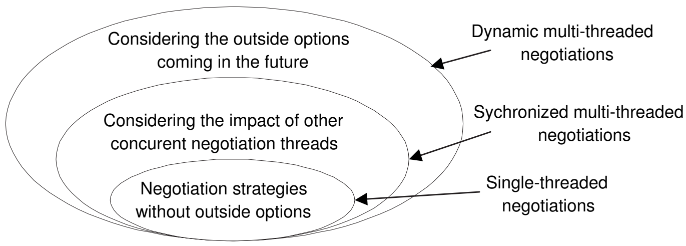

Fig. 1. A nested view of the model

The rest of the paper is organized as follows: Section II presents each specific model. In Section III we provide experimental results. We review related literature in Section IV. Section V concludes.

II. The model

The negotiations follow an *alternating-offers protocol*. In an alternating-offers protocol the negotiators propose and respond alternatively, until one accepts an offer or quits the negotiation, or the negotiation deadline T is reached. The actions of a negotiator at each step in this protocol include: *accept*, *reject and propose an offer*, *quit*. A negotiation has two-sided incomplete information: both negotiation parties do not know the reservation price of each other. Assume the buyer has an estimation of the reservation price of a seller, and the estimation 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. $F(x)$ is identical and independent across all sellers. This probability distribution is called the *prior belief* of the buyer. A negotiation strategy specifies the action at each step conditional on the negotiation history[2], and based on the reservation price and prior beliefs of the negotiators.

There are T periods over the entire horizon for a buyer searching for a negotiation agreement to buy an item from a seller. Let a period be denoted by t, $t = 0, \ldots, T-1$. A buyer needs to reach an agreement with a seller before period T. The potential sellers may come unexpectedly at different times with different reservation prices, and the buyer can negotiate with

[2]The history of a negotiation at time t is a sequence of the two negotiators' actions before t.

the sellers simultaneously. 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. For simplicity we define the value of a seller as the buyer's valuation on the item provided by the seller. Let the value of the seller s_i denoted by v_i. If the buyer reaches an agreement with the seller s_i at x, then the *utility* of the buyer is $v_i - x$.

Denote by $U(D)$ the expected utility of the buyer from a set of negotiation threads D. The *reservation utility* of the buyer in a thread is the lowest utility that the buyer expects from that thread. Given all negotiation threads D, the reservation utility OU_i of the buyer in thread d_i is equal to the expected utility from the outside options: $OU_i = U(D \setminus d_i)$. The *reservation price* of the buyer is the highest price acceptable to the buyer in that thread. The reservation price R_i of the buyer in thread d_i is calculated by $R_i = v_i - OU_i$, i.e., the buyer achieves the reservation utility at the reservation price. If the reservation price in a thread is known, the buyer can apply the single-threaded negotiation model to make the negotiation decisions in that 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 formed by the outside options. In a synchronized multi-threaded negotiation model the outside options at each period for a thread are the other concurrently existing negotiation threads. The synchronized model maps the current outside options to the reservation utility of each thread. The dynamic multi-threaded negotiation model further considers the outside options that may come in the future at uncertain time with uncertain values, and can be viewed as a synchronized model with uncertain threads.

In the following sections II-A, II-B and II-C these models are presented individually.

A. *Single-threaded negotiations*

To calculate the optimal negotiation strategy requires game theoretic analysis of the strategy equilibrium. This analysis is not tractable when both parties have incomplete information and the negotiation is based on an alternating-offers protocol [7] (see Section IV-A for more information on sequential bargaining with two-sided incomplete information). In the AI field some effective heuristic negotiation strategies have been developed to provide formal decision models for automated negotiation agents. Among the generic single-issue quantitative models there are [8], [9], and [1], etc.. Since the focus of this paper is not in designing a single-threaded negotiation strategy, we adopt the time-dependent negotiation strategy that is developed in Faratin et al. [8], for its simplicity, to illustrate the integrative negotiation model with outside options.

In the time-dependent approach time is the predominant factor used to decide which proposal to offer or accept next. For the buyer the proposal to offer or accept is within the interval $[min_b, max_b]$, where max_b is the reservation price of the buyer in the negotiation thread, and min_b is the lower bound of a valid offer (we can reasonably assume min_b=0). Similarly for the seller the proposal to offer or accept is within the interval $[min_s, max_s]$, where min_s is the reservation price of the seller and max_s is the upper bound of a valid offer. Initially a negotiator offers the most favorable value for herself: the buyer starts with min_b and the seller starts with max_s. 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 of time $\alpha_i(t)$, $i \in \{b, s\}$. The proposal x_b^t to be offered by a buyer and the value x_s^t to be offered by the seller at time t, $t \in [0, T-1]$, are as follows:

$$x_b^t = min_b + \alpha_b(t)(max_b - min_b), \quad (1)$$

$$x_s^t = min_s + (1 - \alpha_s(t))(max_s - min_s), \quad (2)$$

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

The time-dependent function $\alpha_i(t)$, $i \in \{b, s\}$, can be defined by a family of polynomial functions[3]:

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

The constant $\beta > 0$ determines the concession pace along time, or the convexity degree of the offer curve as a function of the time (see Figure 2). 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 [2] 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 [10] a negotiator goes to the reservation price rapidly and early. No matter what value β takes, with a constant reservation price, the offer monotonically increases (decreases) with time for a buyer (seller) based on the time-dependent negotiation strategy.

B. *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 and form a synchronized multi-threaded negotiation. The reservation utility that the negotiator should set in one thread d_i is equal to the expected utility from the multi-threaded negotiation formed by all other threads.

[3] Alternatively we can also use the exponential function family, and define $\alpha_i(t) = e^{(1-\frac{t}{T})^{\beta}}$ [8]. 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 [8].

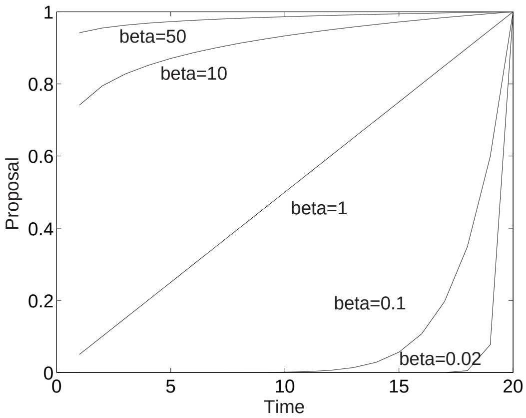

Fig. 2. Offer curves with different β

We present four heuristics to estimate the expected utility from a synchronized multi-threaded negotiation. Three of these heuristics are motivated by the commonality and difference between a synchronized multi-threaded negotiation and a reverse English auction. In a reverse English auction the auctioneer represents the buyer and the bidders are sellers. The buyer wants to buy an item from a single seller. The items from different sellers can be heterogeneous in quality/value. The sellers bid the prices they want to ask for in each step, and the provisional winning bid is the one that brings the highest utility to the buyer (considering the heterogeneity of item values). An outbidden seller decreases her bid from step to step to outperform the winning bid in each step, until she quits or the auction terminates. The auction terminates when no outperforming bid is given, and an agreement is reached at the final outstanding bid with the seller that submits the bid. Since a seller can decrease her bid as low as the reservation price, and the utility of the buyer is equal to the item value minus the price in the agreement, the resulting winner is the seller that has the highest difference between the item value and the reservation price.

Both a synchronized multi-threaded negotiation and a reverse English auction are one-to-many dynamic mechanisms in which sellers compete for one single contract by making sequential concessions. The competition among bidders in a reverse English auction is realized via the winner selection mechanism. In a synchronized multi-threaded negotiation the negotiator achieves indirect competition among opponents by setting a reservation price that is more aggressive than she would without outside options. Facing a more aggressive buyer negotiator the sellers with higher reservation prices or lower quality goods will be less likely to reach an agreement. Although we cannot expect that a synchronized multi-threaded negotiation generates exactly the same outcome as a reverse English auction, the latter, which has been well analyzed theoretically [11], could be used as an approximation of the former mechanism.

However, a synchronized multi-threaded negotiation is also different from a reverse English auction. The major difference lies in the fact that in the negotiation mechanism (with an alternating-offers protocol) both parties can propose and respond, but in an auction only bidders propose and the auctioneer only responds. The auctioneer cannot reject an offer if it is not outbidden. If there is only one bidder, the auctioneer has to accept any valid offer submitted by the bidder, whereas in the negotiation the buyer can always reject an offer and make a counter proposal, even if the opponent is alone. In that sense the buyer in a synchronized multi-threaded negotiation has more negotiation power than the auctioneer in a reverse English auction. To consider this fact, in the approximation of a synchronized multi-threaded negotiation we may follow the reverse English auction with a single threaded bilateral negotiation between the buyer and the auction winner. In the later single-threaded negotiation the buyer can push further the agreement reached in the reverse English auction. Depending on whether or not the single-threaded negotiation is included in the approximation, and how the outcome of a single-threaded negotiation is estimated, we can devise different heuristics of estimating the buyer's expected utility from a synchronized multi-threaded negotiation (see the first three heuristics in Section II-B.2).

Besides these heuristics motivated by auctions, another heuristic is based on the assumption that if any agreement is reached in a multi-threaded negotiation, the agreement is signed with the most *competitive* opponent among all opponents of the threads. Then we approximate the outcome of a multi-threaded negotiation by estimating the outcome of the single-threaded negotiation with the most competitive seller (see the *learning* heuristic in Section II-B.2). 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 of the seller in thread d_i, and $D = \{d_1, \ldots, d_N\}$ is the collection of threads. The amount y_i is the *maximum utility* that the buyer can achieve from the negotiation thread d_i. A higher maximum utility implies that a seller has more space to concede and therefore is more likely to win an agreement with the buyer.

In the rest of Section II-B we shall first derive some relevant variables in part 1), and then present the heuristics for estimating the buyer's expected utility from a synchronized multi-threaded negotiation in part 2).

1) Calculation of relevant variables: We can first have a prediction of a reverse English auction's result [11]. Since a seller gains negative profit at any price agreement lower than her reservation price, a seller has to quit the auction when the current outstanding offer reaches the seller's reservation price, i.e., the seller has to bid lower than her reservation price to outbid that offer. Assume the required minimum bid decrement is infinitely small. In the unique strategy equilibrium a seller continuously decreases her bid until she reaches her reservation price or the auction terminates. The auction terminates when the second most competitive seller asks for her reservation price, and the most competitive seller wins the auction at that price (if we ignore the slight increment of the utility brought by the winner's final offer against the utility by the second most competitive seller's final offer). Therefore the buyer will achieve a utility from the auction equal to the second highest maximum utility.

Since the outcome of the single-threaded negotiation with the most competitive seller depends on how much utility that seller can provide at the upper limit (the width of the "zone of agreement"), we also need to know the distribution of the highest maximum utility.

In the following we derive the probability distributions of the highest and second highest maximum utility of a set of threads $D = \{d_1, ..., d_N\}$. Because the buyer does not know the reservation price of a seller, she does not know the maximum utility of 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 each 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 or equal to y. Let $G^1(y)$ and $G^2(y)$ be the probability distributions of the highest and second highest maximum utility. The probability density functions of $G_i(y)$, $G^1(y)$ and $G^2(y)$ are denoted by $g_i(y)$, $g^1(y)$ and $g^2(y)$ respectively. $G^1(y)$ is equal to the product of the probabilities that the maximum utility is less than or equal to y in each thread. $G^2(y)$ is equal to $G^1(y)$ plus the probability that the highest maximum utility is greater than y, and the second highest maximum utility is less than or equal to y. These probabilities can be calculated by the following formulas:

$$G_i(y) = Pr(v_i - r_i \leq y) = Pr(r_i \geq v_i - y) = 1 - F(v_i - y),$$

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

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

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

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

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

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

2) Estimation heuristics: We provide four heuristic approaches to estimate the expected utility $U(D)$ from a multi-threaded negotiation composed by threads D. To calculate the reservation utility of a thread d in a synchronized multi-threaded negotiation, D is the set of all threads excluding d. Among these heuristics the first three heuristics are based on the approximation with a reverse English auction, while the last one is based on the outcome of a single-threaded negotiation with the most competitive opponent.

- *Conservative estimation:* A synchronized multi-threaded negotiation is approximated by a reverse English auction. The estimated utility of the buyer is equal to the expected second highest maximum utility:

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

  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 $\underline{r}$, and the upper bound of a buyer's valuation is $\bar{v}$, then $\bar{y} = \bar{v} - \underline{r}$.
- *Medium estimation:* In this approach a synchronized multi-threaded negotiation is approximated by a reverse English auction followed by a single-threaded negotiation between the buyer and the winning seller, to consider the bargaining power of the buyer when the buyer can reject and propose in an alternating-offers negotiation. In this approximation the buyer can push further in the single-threaded negotiation the result reached in the auction. Assume the single-threaded negotiation ends at the middle between the buyer's and the winning seller's reservation prices, if the buyer's reservation price is higher than the winning seller's[4]. The buyer's reservation price in the single-threaded negotiation is equal to the price reached in the auction, which brings a utility equal to the second highest maximum utility. At the seller's reservation price in the single threaded negotiation, the buyer gets a utility equal to the highest maximum utility. Therefore the expected utility of the buyer after the single-threaded negotiation is the average of the expected highest and second highest maximum utility:

$$U(D) = (\int_0^{\bar{y}} y g^2(y) dy + \int_0^{\bar{y}} y g^1(y) dy)/2.$$

- *Uniform approximation:* In the medium estimation we assume an agreement can be reached in a single-threaded negotiation as long as there is a zone of agreement, i.e., the buyer's reservation price is higher than the seller's. In this approach we further consider the probability that a negotiation may fail even if there is a zone of agreement, when negotiators do not know each other's reservation prices. Previous research has established an efficient bargaining result with two-sided incomplete information between a buyer and a seller based on game theoretic analysis when both parties' reservation prices follow uniform distributions [12]. Based on this result, if both parties' reservation prices distribute uniformly on $[0, 1]$, an agreement occurs if and only if the buyer's reservation price exceeds the seller's by at least $1/4$. In other words, an agreement cannot be reached if the buyer's reservation price is less than the seller's plus 1/4 of the maximum possible difference between the parties' reservation prices. We can approximate the probability distributions

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

of negotiators' reservation prices by uniform distributions and apply this result to calculate the probability of reaching an agreement. In this heuristic an agreement cannot be reached in the 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 $\bar{y} = \bar{v} - \underline{r}$. In this case of bargaining failure the buyer achieves the reservation utility, which is equal to the second highest maximum utility and is reached in the reverse English auction. If an agreement is reached in the single-threaded negotiation, it is reached at the middle between both parties' reservation prices (same as in the "medium estimation" heuristic). Therefore in the case of reaching an agreement the buyer achieves the medium of the highest and the second highest maximum utility. The probability of reaching an agreement in the single-threaded negotiation is $\int_{\bar{y}/4}^{\bar{y}} g^1(y)dy$, and the buyer's utility is estimated by:

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

- *Learning:* Assume the buyer can learn the distribution of agreements in a single-threaded negotiation based on the previous negotiations [13]. The result of learning is represented by $x(v,r)$, the expected outcome of the negotiation when the reservation prices of the buyer and the seller in the negotiation are v and r respectively. The data on v, r and x can come from the survey of the market, or other third party statistics. This learning is possible in a market where negotiations are repetitive. The matching market in the Navy detailing system is such a market where thousands of sailors are relocated each year and same jobs are repetitively offered on the market. If the seller s_i in the thread d_i is the winning seller, then the probability density of her reservation price r is $f(r)\prod_{d_j \in D\setminus d_i}(1 - F(v_j - v_i + r))$, where the product is the probability that no other thread d_j has the maximum utility $v_j - r_j$ greater than the maximum utility $v_i - r_i$ in thread d_i. The expected utility from the thread d_i is equal to $v_i - x(v_i, r_i)$. Then the expected utility from a multi-threaded negotiation is approximated by the expected utility from the most competitive thread in a single thread situation. The latter utility is equal to the sum of each thread's expected utility $v_i - x(v_i, r_i)$ multiplied by the probability of that thread being the most competitive thread, $\prod_{d_j \in D\setminus d_i}(1 - F(v_j - v_i + r_i))$, conditional on the realization of the seller's reservation price r_i:

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

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 average. The result of learning is expected to be close to the result of a negotiation with $\beta = 1$ for both negotiators:

$$x(v,r) = \begin{cases} \frac{v}{1+v-r} & \text{if } v \geq r \\ \emptyset & \text{otherwise,} \end{cases} \quad (3)$$

assuming the upper bound of an offer is 1 and the lower bound is 0[5].

C. Dynamic multi-threaded negotiations

During an agreement searching process negotiation opponents can be 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. Foreseeing possible arrivals of 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 negotiator can apply the synchronized multi-threaded negotiation model to calculate the appropriate reservation price in each thread. But usually a negotiator is not 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 item value. To set the reservation utility of a thread d, the buyer has to estimate the expected utility from a dynamic multi-threaded negotiation excluding the thread d. In the following of this section we present an approach of estimating the expected utility from a dynamic multi-threaded negotiation.

Following a usual way of modeling uncertain arrivals, we assume the arrival of outside options follows a Poisson process [14]. In each period there is probability p that the negotiator finds an alternative and launches a negotiation thread. The granularity of each period is small enough so that the probability that there are more than one arrival in one period is zero. In a Poisson process the number of arrivals $\eta(\tau, p)$ during an interval with length τ follows a Poisson distribution, $P_{p,\tau}(n) = Pr(\eta(\tau, p) = n) = e^{-p\tau}\frac{(p\tau)^n}{n!}$ [15]. Denote by $\Phi(y) = Pr(v \leq y)$ the probability that an opponent's value is no greater than y. This arrival probability p together with the item value distribution $\Phi(\cdot)$ allows the buyer to forecast the number as well as the quality of the outside options arriving during the rest of the negotiation horizon.

The *state* s_t of the system is defined as the number of past or existing threads n_t[6], and the value of the opponent v_i in each

[5] With $\beta = 1$, the proposal by the buyer at time t is $x_b^t = vt/T$, and by the seller is $x_s^t = 1 - t(1-r)/T$. $x_s^t = x_b^t = v/(1+v-r)$ when $t = T/(1+v-r)$.

[6] We count the past threads in the state because they affect the probability distribution of the maximum utilities of the existing threads. The threads that have survived generally have higher maximum utilities than the threads that have ended earlier.

thread d_i, $s_t = \{n_t, \{v_i\}_{i=1}^{n_t}\}$. The evolution of the system state follows the rule

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

Let $U_t(s_t)$ 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 II-B we can calculate $U(\{n, \{v_i\}_{i=1}^{n}\})$, the expected utility from a synchronized multi-threaded negotiation with n threads and the opponent in thread d_i valued v_i, $i = 1, \ldots, n$. The expected utility in each period follows the recursion

$$U_t(s_t) = (1-p)U_{t+1}(s_t) + pE_v[U_{t+1}(\{n_t+1, \{v_i\}_{i=1}^{n_t} \cup v\})], \quad (4)$$
$$U_{T-1}(s_{T-1}) = U(s_{T-1}).$$

From another perspective $U_t(s_t)$ is the expectation of $U(s_{T-1})$ with respect to s_{T-1}, which depends on s_t. Since the number of arrivals $\eta(\tau, p)$ during an interval with length τ follows a Poisson distribution, equivalently we can calculate the expected utility by:

$$U_t(s_t) = E_\eta[E_{\{v_i\}_{i=n_t+1}^{n_t+\eta}}[U(\{n_t+\eta, \{v_i\}_{i=1}^{n_t} \cup \{v_i\}_{i=n_t+1}^{n_t+\eta}\}]]. \quad (5)$$

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

The expected utility of a dynamic multi-threaded negotiation process at each period with each state can be calculated backward from the final period following Equation 4, or forward following Equation 5 without calculating the the expected utilities in each intermediate period. But even with the forward computing the computation of the expected utility will be expensive because the number of states is exponential with respect to the number of opponent's 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. To simplify the computation we can approximate the result by having the opponent value instances replaced by the expected value $\bar{v}$, i.e.,

$$U_t(s_t) = (1-p)U_{t+1}(s_t) + pU_{t+1}(\{n_t+1, \{v_i\}_{i=1}^{n_t} \cup \bar{v}\})], \quad (6)$$

as an approximation of Equation 4, and

$$U_t(s_t) = E_\eta[U(\{n_t+\eta, \{v_i\}_{i=1}^{n_t} \cup \{\bar{v}\}_{i=n_t+1}^{n_t+\eta}\}]]. \quad (7)$$

as an approximation of Equation 5. 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.

III. Experiments

In Section II 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 different decision models for bilateral negotiations with outside options. In this section we provide experiments to illustrate the models in the solution framework and the performance results. 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 the negotiation strategy. In Section III-A 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 III-B 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.

In the experiments the negotiation deadline $T = 20$. 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 Conceder ($\beta > 1$) or a Boulware ($\beta < 1$). If a negotiator is a Conceder, β^{-1} follows a uniform distribution on $[0, 1]$. If a negotiator is a Boulware, β is a random variable with a uniform distribution on $[0, 1]$. For each arrival probability, we repeat the experiment 100 times and the average utility of the buyer is calculated. The expected utility of a dynamic multi-threaded negotiation process was calculated with the approximation formula, Equation 7 (Section II-C).

A. Reservation utilities and offer curves

We illustrate the impact of outside options on the negotiation strategy by a specific example. In this example $p = 0.2$ and $\beta = 1.262727$[7]. The values and arrival times of outside options in the instance are illustrated in Figure 3.

To illustrate the evolution of the reservation utility of a thread, we collect the reservation utilities of the first thread along time. Figure 4 shows the reservation utilities calculated with different estimation approaches and grouped by the outside option models.

Figure 4 shows that with all different estimation heuristics the reservation utility based on the synchronized model (Section II-B) monotonically increases with time because the number of threads increases with time. It is interesting to note that the reservation utility based on the dynamic model (Section II-C) is not a monotonic function of time. This is because there are two forces that drive the change of the reservation utility: time and arrival of threads. When the

[7] The multiple experiments with different p and β show the same pattern.

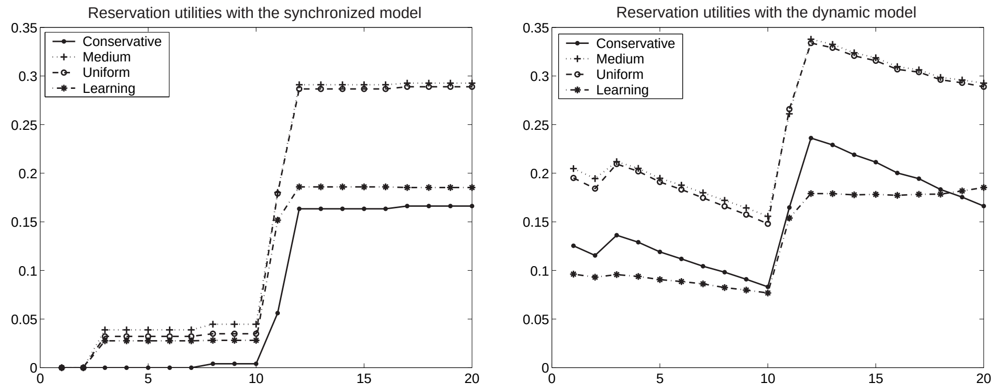

Fig. 4. Reservation utilities

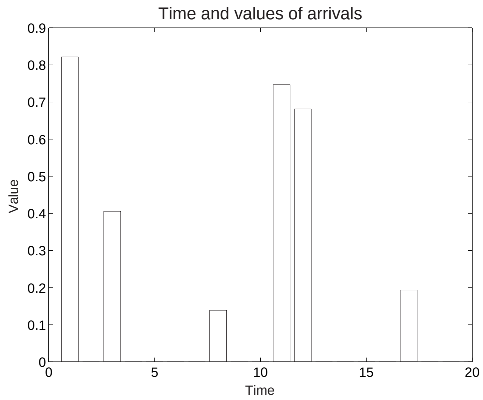

Fig. 3. Time and values of arrivals

negotiator approaches the deadline with time passing, 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. Therefore in the dynamic model the reservation utility deceases with time when there is no new arrival[8]. If there is a new arrival, whether the reservation utility increases or decreases depends on the value of the new opponent. If the value of the new opponent is relatively high, the reservation utility will increase, otherwise decrease. No matter which estimation approach or outside option model is used, the resulting reservation utility with consideration of future outside options is generally higher than without considering the future outside options[8].

The offer curves in the first thread calculated based on different estimation approaches and outside option models are shown in Figure 5. The model noted by "Single" is the model without considering outside options. When the buyer does not consider 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 (e.g., at time 11). This is because the reservation utility of the buyer increases when she sees a new seller that offers a high-value item. When there are no new arrivals, the buyer will concede as time goes by. The concession pace in the synchronized model is the same as in the single-threaded model, but it is greater in the dynamic model. This is because in the dynamic model the buyer expects fewer new arrivals and the reservation utility decreases with passing time. The offers without considering outside options are higher than the offers with considering only concurrent negotiation threads, which are again generally higher than the offers with additional considerations of outside options that may come in the future[8]. This is consistent with the observation that the reservation utility based on the synchronized model is generally less than the one based on the dynamic model.

B. Performance results

In this section we examine and compare the average utilities that a buyer obtains with different models.

Figure 6 is composed of four subplots. Each subplot shows the average utility as a function of the arrival probability based on one reservation utility estimation approach, and with different outside option models[9]. The figure implies that for all estimation approaches and outside option models, the average

[8] An exception is with the learning heuristic. Based on this heuristic the reservation utility in the dynamic model slightly increases with time when no arrivals come after period 17 (see Figure 4), and from period 12 to period 19 the reservation utility (offer) based on the synchronized model is slightly higher (lower) than the one based on the dynamic model (see Figures 4 and 5). This is because in the learning heuristic the winning opponent is assumed to be the opponent with the largest maximum utility $v_i - r_i$. But based on the estimate of the agreement by Equation 3, the expected utility from thread i, if an agreement is reached in this thread, is equal to $v_i - x(v_i, r_i) = v_i(v_i - r_i)/(1 + v_i - r_i)$. It not only depends on the maximum utility $v_i - r_i$, but also on the item value v_i. Based on this equation more threads do not necessarily mean a higher expected utility, although generally it is true. More threads implies that the assumed winning opponent i, the opponent with the highest maximum utility, is guaranteed a higher maximum utility $v_i - r_i$. But the estimated utility from the negotiation with the assumed winning opponent could be lower if the value of this opponent v_i is very low.

[9] The standard deviation ranges from about 70% to 30% of the mean, with the percentage decreasing with the increase of the arrival probability. This high level of standard deviation is due to the introduction of multiple random variables including the number of threads, the negotiation strategy parameter β of a negotiator, the item value in a thread and the reservation price of a seller, that all contribute to the dispersion of the results.

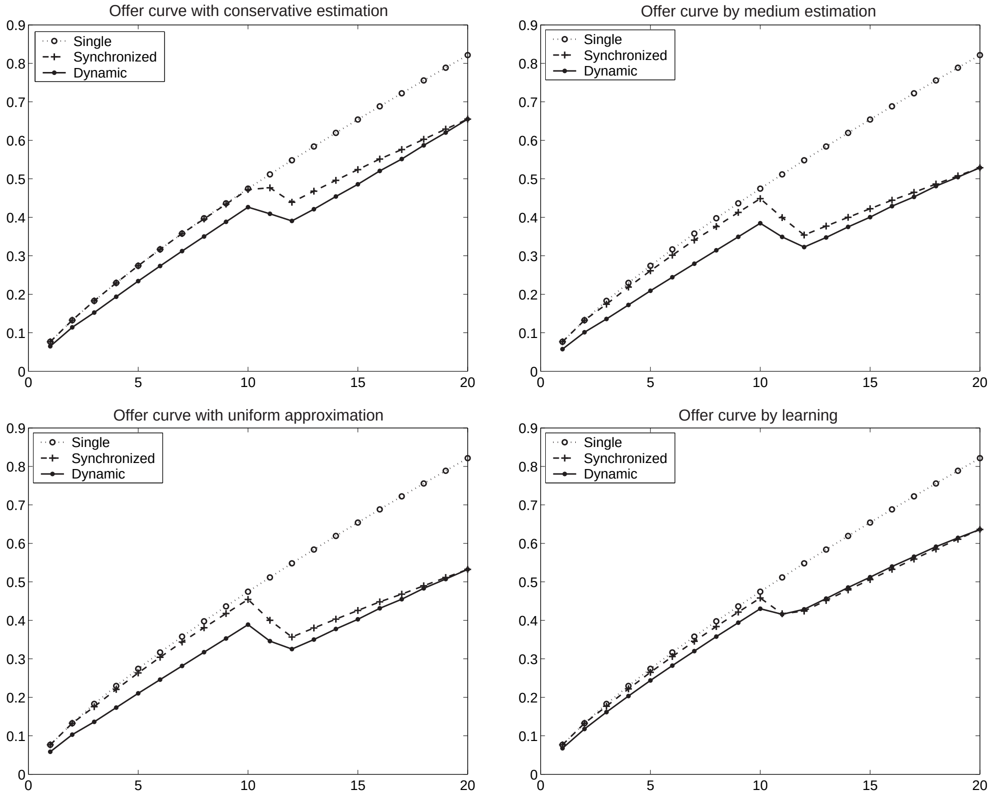

Fig. 5. The offer curves

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. Further more, when the arrival probability is very small, the dynamic model and the synchronized model are close in the average utility, and their advantage over the single-threaded negotiation model is small. When the arrival probability increases, the utility differences between these models also generally increase[10].

Figure 6 also shows that the average utility based on the dynamic model is higher than the one based on the synchronized model, which again brings a higher average utility than the single-threaded model in which no outside option is considered. Compared to the other heuristics, with the learning heuristic the utility difference between the dynamic model and the synchronized model is very small. It is because based on the learning heuristic the reservation utility is less sensitive to the number of outside options. This can also be observed in Figure 4, where the change of the reservation utility with the arrival of outside options is less based on the learning heuristic than based on the other heuristics.

The experimental results show that there is no estimation approach dominating 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.

[10]Again for the learning heuristic this trend of increasing utility difference between the dynamic model and the synchronized model with the increase of the arrival probability is not obvious. This is due to the same reason explained in the footnote 8: with the learning heuristic, more threads do not absolutely bring a higher utility.

IV. Related work

The research work on bilateral negotiations has been conducted in the fields of game theory, and artificial intelligence (AI). The research in game theory focuses on outcomes that satisfy certain axioms, or the strategy equilibrium of agents, based on some rigorous assumptions. Researchers in the field of AI contribute efforts to develop software agents which should be able to negotiate in an intelligent way on behalf of their users. Heuristic approaches are usually used in the complex situations for which game theoretic analysis is untractable. Research in economics and AI have different methodologies and concerns, yet their contributions complement each other. Insights and theoretical foundations developed in game theory provide good heuristics for AI, and the AI approaches provide solutions to negotiations in realistic environments. The computationally feasible solutions provided by AI allow approximate implementations of theoretic results that are developed in a game-theoretic model and that may not be tractable to compute. In the rest of this section we provide a review of related work from both perspectives of game theory and AI. For an extensive version of the review, please refer to Li, et al. [16].

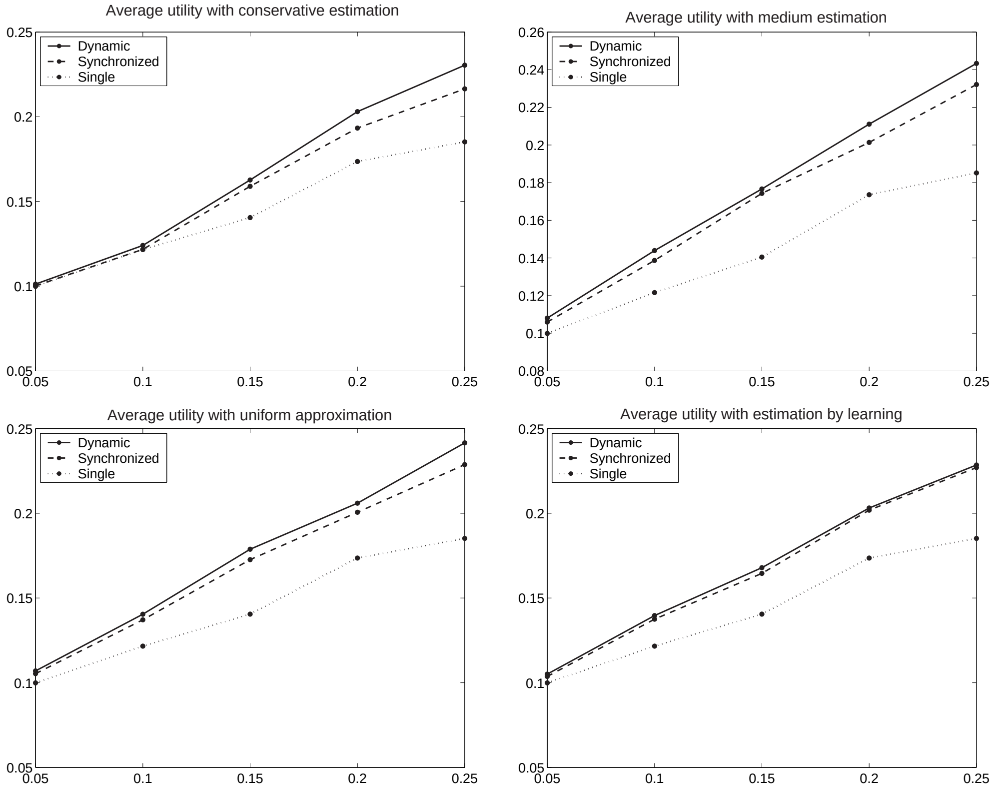

Fig. 6. The average utilities with varying arrival probability

A. The game theory perspective

Game theory can be divided into two branches: cooperative and non-cooperative game theory. *Cooperative game theory* abstracts away from specific rules of a game and is mainly concerned with finding a solution given a set of possible outcomes. The solution is required to satisfy certain plausible properties, such as stability or fairness, which are called *axioms*. *Non-cooperative game theory*, on the other hand, is concerned with specific games with a well-defined set of rules and game strategies, which are known beforehand by the players. A bargaining *strategy* specifies the action of a player at each step given historical information[11] of the negotiation. Non-cooperative game theory uses the notion of an equilibrium strategy to define rational behavior of players, which jointly decide the outcome of a game[17][18].

In cooperative bargaining theory the Nash bargaining solution has been widely used as a modeling tool for negotiations [19]. The Nash bargaining solution is characterized by the payoff pair $s = (x_1, x_2)$ which maximizes the so-called Nash product $(x_1-d_1)^\alpha(x_2-d_2)^\beta$, where d_1 and d_2 are player 1's and player 2's outcomes in case of a disagreement, α and β are the bargaining powers of player 1 and player 2. With outside options, the *disagreement point* (d_1, d_2) can be placed at the *breakdown point* so that d_i is equal to the utility of negotiator i, $i = 1, 2$, from outside options as the negotiation breaks down. The bargaining outcome so generated is called the *split-the-difference* result.

In non-cooperative bargaining theory the outcome of a negotiation depends on the process or protocol of the negotiation. Based on the protocol a negotiation can be a static (one-shot) game or a dynamic (sequential) game. The *alternating-offers protocol* [17] is a widely used protocol for sequential bargaining. In this protocol the two parties take turns to make proposals. In each stage the party who receives a proposal immediately replies by accepting or rejecting the proposal. If she accepts, the negotiation ends; otherwise she makes a proposal to the other party and the negotiation proceeds to the next stage. Rubinstein [20] shows that when utilities are discounted with time, in an infinite bargaining game following the alternating-offers protocol, the unique subgame perfect equilibrium[12] is identical to the Nash bargaining solution, and the equilibrium strategy implies immediate agreement. Based on Rubinstein's model, Binmore et al. [21] established the *outside option principle*. Based on the outside option principle, only those outside options with payoffs that are superior to the payoffs of Rubinstein's equilibrium, or the Nash bargaining solution, have effect on equilibrium strategies. In some cases this principle yields an equilibrium payoff that is different from the split-the-difference outcome. Cunyat [22] re-examined the robustness of the outside option principle based on Rubinstein's bargaining model. His paper argues that the changes that provoke an outside option on a bargaining game depend crucially on if one or both players have the possibility of opting out, and if they can take their outside option either as

[11] There is no historical information if the negotiation is a one-shot game, in which all players take an action simultaneously and then the game ends.

[12] In a *subgame perfect equilibrium* for a sequential game, the strategies constitute a Nash equilibrium for each continuation game.

proposers or as responders. The outside option principle holds only when there is not a gap between the best and the worst continuation subgame payoffs. Muthoo [23] studies a model of the situation in which two players are bargaining face-to-face over the partition of a cake, and one of the players can choose to temporarily leave the negotiating table to search for an outside option. It concludes that the equilibrium outcome does not depend on whether a bargainer is allowed to return to the negotiation table to resume bargaining after having searched for some finite time. Moreover, it shows that the strategic bargaining-search game approximately implements a Nash bargaining solution.

All these papers mentioned above are based on the situation with complete information, i.e., the negotiators know perfectly the preferences and outside options of each other. When a negotiator holds private information on the preference/valuation, it is not feasible to apply the outside option principle. This is because a negotiator does not know the Rubinstein's equilibrium, which requires the knowledge of the preferences on both sides. Hence it is impossible to compare the utility of the Rubinstein's equilibrium with the utility from outside options, even when the latter is known. Our model of the impact of outside options follows the same rule as the split-the-difference principle: the minimum utility acceptable (the disagreement point) is equal to the utility from outside options (the breakdown point).

With incomplete information bargaining inefficiency will be caused by the informational asymmetry: first, an agreement may be delayed, which is necessary for the parties to convey private information credibly, and to screen or signal the private information of the negotiators [24]; second, even when the buyer's valuation exceeds the seller's valuation, the trade occurs with a probability strictly less than one in an equilibrium outcome [12], [7]. This inefficiency in reaching an agreement is reflected in our estimation heuristic, the uniform approximation heuristic. Ausubel and Deneckere [25] establish that the optimal static bargaining mechanism *can* be replicated in sequential bargaining games, in other words, there need not be any additional inefficiency arising from the dynamic nature of bargaining[13]. This conclusion justifies our application of the optimal result of a static bargaining model to sequential negotiations in the uniform approximation heuristic.

Although a unique SPNE can be found by backward inductions for a sequential bargaining game with complete information, the rich information setting in a sequential bargaining game with two-sided incomplete information causes wealth of sequential equilibria, and makes the equilibrium analysis extremely difficult [26], [7]. Actually little is known about the sequential equilibrium strategies with two-sided incomplete information. There are a few papers, such as [27], [28], [29], [30], that characterize, but not explicitly specify, some equilibria in certain restrictive or extreme situations[14]. We adopt a heuristic negotiation strategy to suggest an effective solution when the theoretical equilibrium analysis is intractable.

With outside options specifically considered, Gantner [31] presents a bilateral negotiation model with incomplete information and the alternating-offers protocol. The outside option is modelled as a standard sequential search process. To simplify the analysis the paper assumes that the item has only two values, high or low, for a negotiator. In our paper we consider a general situation in which the item value space is continuous, and outside options are also negotiations that may happen simultaneously and dynamically.

B. *The AI perspective*

Negotiations have received wide attention from the distributed AI community as a pervasive mechanism for distributed conflict resolution between intelligent computational agents [1]. An introduction to negotiation agents is provided in [32]. For a survey on negotiation models in the AI field please refer to Jennings et al. [3] and Gerding et al. [33]. Lomuscio et al. [4] identified the main parameters on which any automated negotiation depends and provided a classification scheme for negotiation models.

The environment that a negotiator is situated in greatly impacts the course of negotiation actions. Instead of focusing on analyzing the strategy equilibrium as a function of (the distribution of) valuations and historical information as in game theory, researchers in AI are interested in designing flexible and sophisticated negotiation agents in complex environments. Faratin et al. [8] devised a negotiation model that defines a range of strategies and tactics for generating proposals based on time, resource, and behaviors of negotiators. We adopt their time-based strategy for the single-threaded negotiation model, but focus on the outside options in the environment.

Sim et al. [34], [35] proposed a market-driven model for designing negotiation agents that make adjustable rates of concession by reacting to some essential market situations that could change over time. The market-driven strategies were further augmented in Sim and Wang [36] with a set of fuzzy rules to enhance the flexibility of negotiation agents. The market situations include trading opportunities, competition, remaining trading time, and eagerness. Multiple trading opportunities in the market can be regarded as outside options against each other for a negotiator. In their model the number of trading opportunities influences the aggregated probability of conflict, which determines the probability of completing a deal in the current negotiation cycle. With more trading opportunities, the probability of completing a deal is higher, and it follows that the negotiator's concession is smaller in the next cycle based on the spread decision function.

Krovi et al. [37] devised a genetic algorithm-based model of negotiations, and examined the impact of task, agent, and communication characteristics on agent behaviors as well as the outcome of negotiations. In our work the outside options affect the negotiation strategies via their impact on the reservation price. An agent is not only able to adjust the strategy reactively to the emerging outside options (in the synchronized model), but also set the reservation utility proactively by predicting the arrival and impact of future outside options (in the dynamic

[13] This is based on the assumption that the valuation distribution functions exhibit monotonic hazard rates.

[14] For example, the preferences of negotiators are uniformly distributed [27], delay of communications is allowed and there is no negotiation deadline [29], [27], or one party makes all the offers [27], [28], [30].

model). Since the objective of a negotiator is to maximize the expected utility achieved in the entire negotiation horizon, acting proactively can prevent some short-sighted behaviors that do not take into account the impact of the behavior on the future. For example, a negotiator would over compromise and increase the probability of reaching less attractive deals without foreseeing the possible outside options arriving in the future. Or a negotiator would be too tough to reach a deal without predicting the reducing probability of having more outside options with time proceeding. The benefit of appropriately planning the future is shown by the advantage of the dynamic model against the synchronized model in the experimental results.

Among the papers that specifically and particularly consider outside options, there are Nguyen and Jennings [38], [39], and Byde and Preist et al. [40], [41]. Nguyen and Jennings [38], [39] presented a heuristic model that enables an agent to participate in multiple, concurrent bi-lateral negotiations without considering the future arrivals of outside options. In this model the buyer that has outside options can accept an offer from a seller, with the agreement binding only on the seller but not on the buyer. In other words, the buyer can decline the agreement that is not finalized if she finds a better deal later. This protocol is extremely buyer-biased as the buyer is guaranteed the best offer she can find from all different threads. In reality the buyer is usually not a monopoly player. Sellers may also have outside options and can opt out or withdraw an offer before the buyer finalizes the decision. The alternating offers protocol adopted in our paper is fair to both sides. With this protocol a negotiator on each side has to respond to an offer immediately and the response binds on both parties. This obligation on binding also motivates both parties to be more serious on making offers and improves the market efficiency.

Priest et al. [41] developed algorithms for agents to participate in multiple parallel English auctions for the purchase of similar goods. Byde et a. [40] further developed a decision theoretic framework and a heuristic algorithm that enable an agent to make decisions about purchasing multiple goods from multiple auctions that operate different protocols. Although in their theoretic framework the multiple auctions can open in different time, in the heuristics it is assumed that all auctions proceed in synchronized rounds. In our work we consider multiple bilateral negotiations instead of auctions.

V. Conclusion

In this paper we provide a negotiation decision model that considers outside options. This model is motivated by real world situations and proposed for solving negotiation problems in the Navy detailing process. We have presented a computational and analytic model to model outside options, and a solution that integrates the outside option models and negotiation strategies. Experimental results show that this model provide higher utility to a negotiator than the models that do not consider outside options appropriately.

We have proposed several heuristic approaches based on game theory and probability theory to set appropriate reservation prices considering outside options, and built the solution based on the time-dependent negotiation strategy. We do not claim that the heuristics or the negotiation strategy we provide in this paper 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.

In this 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 in this solution framework if the prior belief on the opponent's reservation price also includes the probabilistic information on her outside options. Alternatively and more explicitly, we can introduce the probability of a seller's opting-out action to reflect the availability of outside options to the seller. As a buyer has no information about the available outside options of a seller, she sees the opting-out of a seller as a random event, and the opting-out probability depends on time and the current offer. Learning/modeling this opting-out probability and incorporating the probability in the negotiation strategy are augmentations for future work.

Acknowledgement

This work was supported in part by the Office of Naval Research, 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.

The authors thank the anonymous referees for their valuable comments and suggestions.

References

[1] J. Rosenschein and G. Zlotkin, *Rules of Encounter*. MIT Press, 1994.

[2] H. Raiffa, *The Art and Science of Negotiation*. Harvard University Press, 1982.

[3] N. R. Jennings, P. Faratin, A. R. Lomuscio, S. Parsons, C. Sierra, and M. Wooldridge, "Automated negotiation: prospects, methods and challenges," *Int. J. of Group Decision and Negotiation*, vol. 10, no. 2, pp. 199–215, 2001.

[4] A. R. Lomuscio, M. Wooldridge, and N. R. Jennings, "A classification scheme for negotiation in electronic commerce," *International Journal of Group Decision and Negotiation*, vol. 12, no. 1, pp. 31–56, 2003.

[5] D. Gale, *Strategic foundations of general equilibrium: dynamic matching and bargaining games*. Cambridge University Press, 2000.

[6] K. Sycara, "Utility theory in conflict resolution," *Annals of Operations Research*, vol. 12, pp. 65–84, 1988.

[7] L. Ausubel, P. Cramton, and R. Deneckere, *Bargaining with incomplete information*. Elsevier Science, 2002, vol. 3, ch. 50, robert J. Aumann and Sergiu Hart, eds.

[8] P. Faratin, C. Sierra, and N. R. Jennings, "Negotiation decision functions for autonomous agents," *International Journal of Robotics and Autonomous Systems*, vol. 24, no. 3-4, pp. 159–182, 1998.

[9] D. D. Zeng and K. Sycara, "Bayesian learning in negotiation," *International Journal of Human-Computer Studies*, vol. 48, pp. 125–141, 1998.

[10] D. G. Pruitt, *Negotiation Behavior*. Academic Press, 1981.

[11] V. Krishna, *Auction Theory*. Academic Press, 2002.

[12] R. B. Myerson and M. A. Satterthwaite, "Efficient mechanisms for bilateral trading," *Journal of Economic Theory*, vol. 29, pp. 265–281, 1983.

[13] K. Sycara, "Machine learning for intelligent support of conflict resolution," *Decision Support Systems*, vol. 10, pp. 121–136, 1993.

[14] S. Lippman and J. McCall, *The economics of uncertainty: Selected topics and probabilistic methods*. Amsterdam: North-Holland Publishing Company, 1981, vol. I, ch. 6, In K.J. Arrow and M. Intriligator, editors.

[15] S. M. Ross, *Introduction to Probability Models*. Academic Press, 2000.

[16] C. Li, J. Giampapa, and K. Sycara, "A review of research literature on bilateral negotiations," Carnegie Mellon University, Tech. Rep. CMU-RI-TR-03-41, 2003, URL: http://www-2.cs.cmu.edu/∼softagents/papers/negotiation_lit.ps.

[17] S. Napel, *Bilateral Bargaining: Theory and Applications*, ser. Lecture notes in economics and mathematical systems. Springer, 2002.

[18] M. J. Osborne and A. Rubinstein, *A Course in Game Theory*. MIT Press, 1994.

[19] J. F. Nash, "The bargaining problem," *Econometrica*, vol. 28, pp. 155–162, 1950.

[20] A. Rubinstein, "Perfect equilibrium in a bargaining model," *Econometrica*, vol. 50, no. 1, pp. 97–109, January 1982.

[21] K. Binmore, A. Shaked, and J. Sutton, "An outside option experiment," *Quarterly Journal of Economics*, vol. 104, no. 4, pp. 753–770, 1989.

[22] A. Cunyat, "Strategic negotiations and outside options," Department of Economic Analysis, University of Valencia, Tech. Rep., 1998.

[23] A. Muthoo, "On the strategic role of outside options in bilateral bargaining," *Operations Research*, vol. 43, no. 2, pp. 292–297, 1995.

[24] J. Kennan and R. Wilson, "Bargaining with private information," *Journal of Economic Literature*, vol. 31, no. 1, pp. 45–104, 1993.

[25] L. Ausubel and R. Deneckere, "Efficient sequential bargaining," *Review of Economic Studies*, vol. 60, pp. 435–462, 1993.

[26] D. Fudenberg and J. Tirole, *Game Theory*. MIT Press, 1991.

[27] P. Cramton, "Bargaining with incomplete information: an infinite-horizon model with continuous uncertainty," *Review of Economic Studeis*, vol. 51, pp. 579–593, 1984.

[28] I.-K. Cho, "Uncertainty and delay in bargaining," *Review of Economic Studies*, vol. 57, pp. 575–596, 1990.

[29] P. Cramton, "Strategic delay in bargaining with two-sided uncertainty," *Reivew of Economic Studies*, vol. 59, pp. 205–225, 1992.

[30] L. Ausubel and R. Deneckere, "Durable goods monopoly with incomplete information," *Review of Economic Studies*, vol. 59, pp. 795–812, 1992.

[31] A. Gantner, "Bargaining, search and outside options," Department of Economics, University of California, Santa Barbara, Tech. Rep., 2002.

[32] M. Wooldridge, *An Introduction to MultiAgent Systems*. John Wiley and Sons, 2002.

[33] E. H. Gerding, D. van Bragt, and J. L. Poutré, "Scientific approaches and techniques for negotiation: a game theoretic and artificial intelligence perspective," CWI, Tech. Rep. SEN-R0005, 2000.

[34] K. M. Sim, "A market-driven model for designing negotiation agents," *Computational Intelligence*, vol. 18, no. 4, pp. 618–637, 2002, special issue in Agent Technology for E-Commerce.

[35] K. M. Sim and C. Y. Choi, "Agents that react to changing market situations," *IEEE Transaction on Systems, Man and Cybernetics, Part B: Cybernetics*, vol. 33, no. 2, pp. 188–201, April 2003.

[36] K. M. Sim and S. Y. Wang, "Flexible negotiation agent with relaxed decision rules," *IEEE Transaction on Systems, Man and Cybernetics, Part B: Cybernetics*, vol. 34, no. 3, pp. 1602–1608, June 2004.

[37] R. Krovi, A. C. Graesser, and W. E. Pracht, "Agent behaviors in virtual negotiation environments," *IEEE Transaction on Systems, Man and Cybernetics, Part C: Applications and Review*, vol. 29, no. 1, pp. 15–25, February 1999.

[38] T. D. Nguyen and N. Jennings, "A heuristic model for concurrent bilateral negotiations in incomplete information settings," in *Proceedings of International Joint Conferences on Artificial Intelligence*, Mexico, 2003, pp. 1467–1469.

[39] ——, "Concurrent bi-lateral negotiation in agent systems," in *14th International Workshop on Database and Expert Systems Applications (DEXA'03)*, Prague, Czech Republic, September 2003.

[40] A. Byde, C. Preist, and N. R. Jennings, "Decision procedures for multiple auctions," *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-agent Systems (AAMAS'02)*, 2002.

[41] C. Preist, A. Byde, and C. Bartolini, "Economic dynamics of agents in multiple auctions," *Proceedings of the 4th International Conference on Autonomous Agents (AGENTS'01)*, 2001.