

Modelling Bilateral Negotiations in the Navy
Detailing Process

A
Carnegie Mellon University
Robotics Institute
Technical Report
CMU-RI-TR-03-47

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

November 16, 2003

Abstract

We propose 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. 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. The modules are described individually, and extensions are discussed for their application in more general situations.

1 Introduction

In this report we develop a model of negotiations in the Navy detailing process. The selection of the model is based on the following principles:

- High fidelity: The model should reflect the important features and concerns of negotiations in the Navy detailing system. The essential issues include: the outside options, uncertainty and incomplete information, etc. Some simplifications and assumptions, however, are necessary to make the quantitative description and analysis possible without sacrificing the fidelity of the model.
- Modular: We want the model to be modular so that one module encapsulates certain relatively independent functions. Changes in the internal process of a module should not require rebuilding the whole model.
- Extensible: We want to identify the key interfaces for the model to integrate other functions that are not currently included, and to accommodate alternative considerations.

Bilateral negotiations is an important mechanism to implement flexible and distributed matching in the Navy detailing system [5, 10]. To build a quantitative model to characterize the negotiation problem is necessary for providing rational decision support and building automated negotiation systems. A negotiation strategy is a mapping from input information about the *environment* to a sequence of decisions. By environment we mean all factors that impact the negotiation outcome, for example, valuation of the matching opportunity (i.e. suitability of a posted job to a sailor's experience and preferences), the possible valuation held by the "opponent"¹, and the deadline for reaching an agreement. In the literature review [10] we have provided an extensive survey on research work in negotiations in both fields of economics and artificial intelligence. The previous work on negotiation decisions has been focused on the design of negotiation strategies assuming all inputs are given and static. But in a realistic situation such as the Navy detailing process it is not straightforward to acquire the input information because of a complex, uncertain and dynamic environment. How to set the input depends on how the decision maker understands the environment and interprets the impact of the environment on the negotiation.

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

¹Here we refer to the other party against whom a party is negotiating as the "opponent". So, if a sailor must make a negotiating decision, the "opponent" is a command and vice versa.

for the job, and a sailor can be informed of more than one job vacancies that interest him. These alternatives are called *outside options*, and they contribute to the environment of one negotiation. 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.

In the Navy detailing process, outside options can exist *concurrently* with a negotiation, or come *sequentially* in the future. 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 all 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 impact the input to a negotiation decision model as a part of the environment. Existence of outside options changes the utility that an agent expects from the current negotiation, and hence the agreement that is acceptable for the agent in the current negotiation [18]. From the above modelling considerations it is clear that the Navy detailing process is complex and cannot be handled by existing negotiation models.

Based on the above modelling principles, we propose a model that consists of three nested modules: single-threaded negotiations, synchronized multi-threaded negotiations and dynamic multi-threaded negotiations. The model of *single-threaded negotiations* can be viewed as the core negotiation decision function that contains all information on a given environment. The model of *synchronized multi-threaded negotiations*

builds on the single-threaded negotiation model and considers the presence of concurrently available outside options. The model of *dynamic multi-threaded negotiations* expands the synchronized multi-threaded negotiation model by considering the uncertain outside options that could arrive dynamically in the future. This integrated model provides a cohesive framework to build effective negotiation strategies with outside options by considering different decision factors separately and additively.

The rest of the report is organized as follows: Section 2 presents an integrative view of the modelling framework, and the specific model of each module. In Section 3 we discuss how to extend the model to some more general situations. Section 4 concludes.

2 The model

To model the system we have to first intuitively understand how outside options impact the negotiation strategies. We claim that *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. Call the utility of the worst acceptable agreement in a negotiation the *reservation utility* of the negotiation. The reservation utility is different from the reservation price if the value/cost of the negotiated good is not zero. The reservation price can be defined as the value minus the reservation utility for the buyer. It is the cost plus the reservation utility for the seller.

The reservation utility in a negotiation is equal to the utility that the negotiator can get without an agreement with the opponent, and it depends on the availability of outside options. If there are no outside options now or in the future, then the current negotiation is the only matching opportunity that the negotiator can have. If the negotiator breaks up with the opponent, she gets zero utility. Therefore the reservation utility in the negotiation is zero. Based on the reservation utility the negotiator decides the reservation price in the negotiation. For a command that negotiates the incentive pay with a sailor, the utility of the command from a matching agreement with the sailor can be defined as the value of the sailor to the command minus the incentive pay that is offered in the agreement. Hence in the situation without outside options, the reservation price held by a command is equal to the sailor's value. Likewise for the

sailor the reservation price, or the minimum acceptable incentive pay, is equal to how much he values the loss (for example, suffering of the personal life) by taking the job.

If there are outside options, then the reservation utility in a negotiation is equal to the *expected* utility from outside options that are available now or later. The utility from outside options is measured on expectation because of the uncertainty of the availability and quality of outside options. We should remark that the claim is based on the assumption that a thread is *non-resumable* [10]. In a resumable thread a negotiator can leave the negotiation table temporarily for discovering more information in other negotiation threads, and come back to resume the negotiation if necessary. For a resumable thread the outside option also includes the option to come back to the negotiation table, which influences the negotiation strategy. But the value of this option also depends on the negotiation strategy before the negotiator opts out, and the decision of when to leave and return to the table. This recursive influence between this option and the negotiation strategy makes the decision untractable to analyze. In this work we assume that the negotiation threads are non-resumable [10]. This assumption is not restrictive because prior research on bargaining [4, 13] shows that having the option to suspend and resume a negotiation does not improve the utility of a negotiator in typical bargaining situations. Also, availability of outside options to either party reduces the likelihood that the waiting party will remain committed to a suspended thread.

Design of an effective negotiation strategy can be divided into two parts: the first is the design of 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. 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 we can assume there are no outside options coming in the future. The outside options are those negotiation threads that concurrently exist with the thread under consideration. In other words, all negotiation threads are assumed to start at the same time. Therefore there is no uncertainty about outside options in terms of both the thread number and opponents. We call this model on negotiations with only concurrently available outside options *synchronized multi-threaded negotiations*. On the second level we also consider the outside options that may come dynamically in the future. Hence the number of threads that the negotiator would be involved in is a random variable and changes with time. The opponents in the future threads are also uncertain. 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 on the number of 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. Table 1 summarizes the essential additional inputs of each module from the module on the last complexity level.

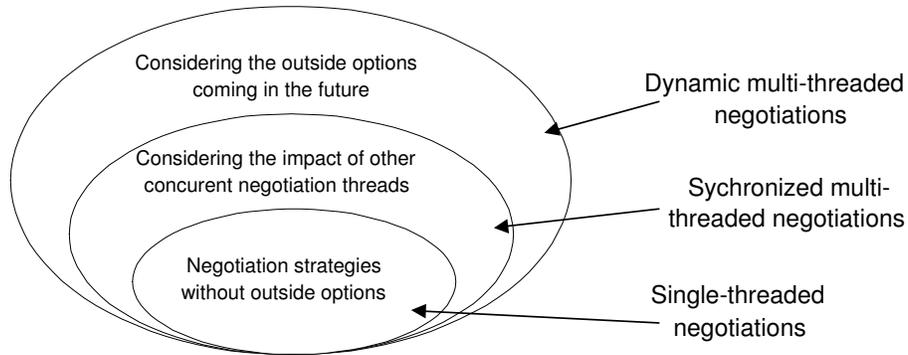


Figure 1: A nested view of the model

Modules	Additional inputs
Single-threaded	Reservation price, Prob. dist. of opponents' reservation prices
Synchronized multi-threaded	Number of threads, Opponents' values
Dynamic multi-threaded	Arriving prob. of outside options, Prob. dist. of opponents' values

Table 1: Additional inputs of each module

In the following sections 2.1, 2.2 and 2.3 these models are presented individually.

2.1 Single-threaded negotiations

The single-threaded negotiation model determines the negotiation strategy of a negotiator given necessary inputs. The nature of the necessary inputs depends on the information structure and the negotiation protocol. The negotiations in the Navy detailing process are two-sided incomplete information situations, because a command

does not know a sailor’s personal preferences for incentive features, and the limit to which a command can offer incentive features may depend on the specific situation and budget of individual commands. 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, 15]. The negotiation strategy based on an alternating-offers protocol specifies the decisions for both proposal generation and response to an offer².

In an alternating-offer negotiation with two-sided incomplete information, the necessary inputs are the prior belief on the opponent’s *type*, the negotiator’s reservation price and the time deadline of the negotiation. The *type* of the opponent is a game theoretic notion that includes the private information of the opponent that affects her negotiation strategy and thus the outcome of the negotiation. For example, the type can be how the negotiator evaluates the possible agreements, the attitude towards risk, time preferences, etc. In the real world there can be many different pieces of information about a negotiator that are hidden from the opponent. It is not possible to include all possible private information in the model because that would make the model intractable. We need to choose the private information that is the most important factor in the negotiation decision. We model a negotiator’s type by the reservation price. Compared to other private information such as risk attitude, reservation price varies much more with different opponents. The negotiation outcome is very sensitive to the change of the opponent’s reservation price, which determines what she will likely offer or accept.

The risk attitude can be indirectly reflected in the reservation price. An agent in the role of a seller usually has a lower reservation price if she is *risk-averse*, or a higher one if she is a *risk-lover*, than a *risk-neutral* agent. A negotiator is *risk-neutral* if she is indifferent between a deterministic situation and a random situation which brings the same utility *on expectation* as the utility from the deterministic situation. If a risk neutral negotiator has to decide either to accept an offer or reject it, the decision only depends on the comparison of the expected utility from continuing the negotiation to the deterministic utility of the current offer. She will take the current offer if and only if the former utility is less than the latter. The uncertainty of the actual utility, or

²The alternating-offers strategy can be customized, possibly with some simple changes that depend on the strategy, to ultimatum negotiations by taking the corresponding decision part, generating a proposal, or responding to an offer.

risk, in the stochastic situation does not impact the decision. A risk neutral agent is indifferent to either a stochastic situation or to a deterministic situation, if they bring the same utility either on expectation or with certainty. A *risk-averse* agent prefers more the deterministic situation considering the probability that the actual utility may be lower than the deterministic utility, while a *risk-lover* prefers the stochastic situation considering the probability of higher actual utilities. In this model we assume negotiators are risk-neutral. We will discuss how to extend the model to the situation where a negotiator may be risk-averse or a risk-lover. A command or a sailor has no special preferences with respect to time, because their goal is to reach a desirable agreement before the end of the search or application period. There is no difference on utility between getting an agreement today or tomorrow, as long as it is earlier than the deadline.

Let $i \in \{a, b\}$ represent negotiators in a single thread. Negotiator a is the one who prefers a higher value on the negotiated issue and negotiator b is the one who prefers a lower value. The negotiated issue can be the incentive pay, the arriving time, or the vacation that concerns a command and a sailor³. Let's assume the negotiated issue is the incentive pay. Negotiator a is the sailor and negotiator b is the command, since a sailor typically prefers higher incentive pay whereas a command prefers lower. The reservation price of negotiator i is denoted by R_i , $i \in \{a, b\}$. For the sailor, or negotiator a , R_a is equal to the minimum incentive pay that he can accept. For the command, or negotiator b , R_b is the maximum incentive pay that he is willing to offer to the sailor.

Each negotiator knows her own type, but not the type of the opponent. However, a negotiator can have some estimation of the opponent's type, 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 positions 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 uncertain types is characterized by a probability distribution of the types. Negotiator a believes that the probability of R_b being less than x is $F_a(x)$. This is called the *prior belief* of negotiator a on the type of negotiator b . Likewise the prior belief of negotiator b is $F_b(x)$.

Let the deadline of the negotiation be T , and the negotiation starts at time 0. The history H^t of a negotiation at time t , $t \geq 0$, 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 . In a negotiation following an alternating-offers protocol, the negotiators propose and re-

³In this work we focus on single-attribute negotiations. Extension of this model to multi-attribute negotiations is discussed at the end of this report.

spond alternatively, until one accepts an offer or quits the negotiation⁴. Therefore the history of an alternating-offers negotiation at time t is a sequence of proposals, i.e., $H_t = \{x_b^0, x_a^1, x_b^2, x_a^3, x_b^4, \dots, x_a^t(x_b^t)\}$, where x_i^m is the proposal submitted by negotiator i at time t . The 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 $A_i^t \in \{\text{accept}, \text{reject and propose } x_i^{t+1}, \text{quit}\}$. Some previous work, such as Zeng and Sycara [22], Faratin et al. [2], and Huang and Sycara [7], provides valuable references on designing effective single-threaded negotiation strategies.

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.

Let the number of threads be N , and the collection of threads be $D = \{1, 2, \dots, N\}$. Denote the collection of threads other than thread d be $D \setminus d$. If the expected utility from the multi-threaded negotiation process formed by $D \setminus d$ is OU_d , then the reservation price in thread d is $R_d = v_d - OU_d$, where v_d is the value of having the job filled by the opponent in thread d if the negotiator is a command, or the value⁵ of the job if the negotiator is a sailor. Note that the reservation price can be positive or negative. A positive price means the negotiator pays, while a negative price means the negotiator is paid. It is possible that if a sailor prefers a job very much, he may even be willing to give up some standard benefit, in other words, to pay the command, if it is allowed. Given R_d the negotiator can negotiate in thread d following a strategy derived from the single-threaded negotiation model. The problem is how to compute the reservation utility OU_d . As has been noted OU_d is the expected utility from a synchronized multi-threaded negotiation process $D \setminus d$.

⁴One could quit the negotiation because an agreement is reached in another negotiation thread, or a proposal at her reservation price is rejected by the other party.

⁵The value of a job to a sailor reflects how much the sailor prefers the job. The value can be negative if the job requires sacrifice of some personal preferences, for example, separation from the family or long time stay on the sea.

We can approximate the expected utility from a multi-threaded negotiation by viewing it as a *pseudo* auction mechanism. We call it a pseudo auction mechanism because the negotiator can induce competition between opponents in different threads as in an auction, but the resulting price is different. The negotiator can reveal to an opponent the existence of other threads and the current best proposal in other threads, and hence stimulate competition among opponents in different threads. When the current best proposal across the threads reaches the reservation price of an opponent, the opponent has to quit because she cannot outbid that best proposal. Stimulating competition lowers the utility expectation of opponents, and thus increases the bargaining power of the negotiator in each thread [20]. The negotiation terminates either with a single-threaded negotiation or with a proposal of the negotiator being accepted by an opponent. Generally the opponent who is the last one left or is the first one to accept a proposal from the negotiator is the most competitive one among all opponents. Therefore a multi-threaded negotiation provides a process to discover the winner, which is similar to an auction. The difference between a multi-threaded negotiation process and an auction includes: (1) the communications in a multi-threaded negotiation are not synchronized across the threads as they are in an auction; and (2) in a multi-threaded negotiation the deal between the negotiator and the winning opponent may need to be determined by continued bargaining after all other threads are terminated, but in an auction the deal is concluded at the same time when the winner is discovered.

Estimation of the expected utility from a multi-threaded negotiation can be divided into two parts: the first part is to estimate the type of the winner, and the second part is to approximate the expected utility from a single-threaded negotiation with the winner. In the following we describe the approaches for these two parts separately.

Let the reservation price of the opponent in thread d be r_d . The *maximum utility* u_d of an opponent d is the utility that the negotiator can get from an agreement with the opponent at the opponent's reservation price, i.e., $u_d = v_d - r_d$, where v_d is the opponent's value. Then the probability of the maximum utility being less than y is $G_d(y) = Pr(v_d - r_d \leq y) = Pr(r_d \geq v_d - y) = F(v_d - y)$, where $F(\cdot)$ is the prior belief of the negotiator on opponents' reservation prices. Maximum utility is the indicator of the competitiveness of opponents. The winning opponent has the highest maximum utility among the opponents. Let $G^1(y)$ be the probability distribution of the winning opponent's maximum utility. Then $G^1(y) = \prod_{d \in D} G_d(y)$. Similarly we can construct the probability distribution of the second highest maximum utility among the opponents: $G^2(y) = G^1(y) + \sum_{d=1}^N (1 - G_d(y)) \prod_{e \in D \setminus d} G_e(y)$. Assume the opponent in thread d is the winner. Then the probability distribution of the winner's reservation price is $F^1(x) = Pr(r_d \leq x) = Pr(u_d \geq v_d - x) = 1 - G^1(v_d - x)$, and it is the prior belief that the negotiator holds on the winning opponent's type in the single-threaded

negotiation with the winner. Note that the prior belief on the winning opponent $F^1(\cdot)$ is different from the one on a normal opponent $F(\cdot)$.

The next question is how to estimate the expected utility U from a single-threaded negotiation given the reservation price and the prior belief. The reservation price of the negotiator in the final single-threaded negotiation is v_d if the winning opponent is the one in thread d . The prior belief on the opponent's reservation price is $F^1(\cdot)$. Theoretically the expected utility should be calculated based on a game-theoretic strategy equilibrium of the negotiation [16, 15]. But for a alternating-offers negotiation with two-sided incomplete information the equilibrium is not tractable to compute. We suggest the following heuristic approaches to make the estimation.

- **Conservative estimation:** One conservative way is to estimate that the negotiator achieves the utility which is equal to the maximum utility of the second most competitive opponent. The winning opponent is discovered at the moment when the second most competitive opponent is rejected. At that moment the second most competitive opponent proposes an offer at her reservation price, which gives up all her utility to the negotiator, but still cannot outbid the most competitive opponent. If the recession of the winning opponent in the continued single-threaded bargaining is ignored, the negotiation ends up with an agreement with the winning opponent slightly better than the last proposal given by the second most competitive opponent. In that situation the utility of the negotiator is equal to the expected second highest maximum utility among the opponents. This is a conservative estimation because the negotiator can get higher utility by continuing bargaining with the winning opponent. The expected second highest maximum utility can be calculated based on the probability distribution $G^2(\cdot)$.
- **Medium estimation:** Assume the bargaining ends at the middle point between both parties' reservation prices if the command's reservation price is higher than the sailor's⁶. 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 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 [14, 3]. Based on this

⁶If the command's reservation price is lower than the sailor's, there is no "zone of agreement" and the negotiation will fail.

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 $[0, 1]$. We can approximate the probability distributions of negotiators’ types by uniform distributions and apply this result to calculate the probability of reaching an agreement. When an agreement is reached, it is reasonable to assume that it is at the middle point between both parties’ reservation prices.

- Learning: Learn the probability of reaching an agreement and the distribution of agreements based on the previous negotiations [21].

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 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 [11, 12, 17]. There are T periods over the entire horizon of a detailing window. Let a period be denoted by t , $t = 0, \dots, T - 1$. In each period there is probability p that the negotiator 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. The value v of an opponent follows the probability distribution $H(y) = Pr(v \leq y)$, where $H(y)$ is the probability that an opponent’s value is no greater than y . The reservation price r of an opponent follows the prior belief $F(x) = Pr(r \leq x)$. 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 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 $U_t(s_t)$ 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 2.2 we can calculate $U(\{n, \{v_d\}_{d=1}^n\})$, 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

$$U_t(s_t) = (1 - p)U_{t+1}(s_t) + pE_v[U_{t+1}(\{n_t + 1, \{v_d\}_{d=1}^{n_t} \cup v\})], \quad (1)$$

$$U_{T-1}(s_{T-1}) = U(s_{T-1}).$$

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

$$U_t(s_t) = E_\eta[E_{\{v_d\}_{d=n_t+1}^{n_t+\eta}} [U(\{n_t + t, \{v_d\}_{d=1}^{n_t} \cup \{v_d\}_{d=n_t+1}^{n_t+\eta}\})]] \quad (2)$$

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

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 changes 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 1 or forward following Equation 2. The computation will be very heavy 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.,

$$U_t(s_t) = (1 - p)U_{t+1}(s_t) + pU_{t+1}(\{n_t + 1, \{v_d\}_{d=1}^{n_t} \cup \bar{v}\}) \quad (3)$$

$$U_t(s_t) = E_\eta[U(\{n_t + t, \{v_d\}_{d=1}^{n_t} \cup \{\bar{v}\}_{d=n_t+1}^{n_t+\eta}\})] \quad (4)$$

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 Discussions

Our focus in this work is to consider negotiation strategies when negotiators face uncertain and dynamic outside options. To simplify the presentation and make the analysis tractable, we have restricted the models to single-attribute negotiations in which negotiators are risk-neutral. In the Navy detailing process a sailor and a command may negotiate more than one incentive features, i.e., the negotiations may be multi-attribute negotiations. A negotiator may be risk-averse or a risk-lover. In the following sections we discuss how to extend the model in this work to the more general situations where a negotiation involves multiple attributes, or negotiators have different risk attitudes.

3.1 Multi-attribute negotiations

The modelling framework presented in this report can be extended to the situation of negotiating for multiple attributes, although we have based the work on single attribute negotiations. According to the form of the utility function, a multi-attribute negotiation can be “competitive” or “integrative” [8, 19, 9]. The negotiation is competitive if negotiators use an additive scoring function of the attribute values to evaluate the agreements. In this situation the utility brought by one attribute does not depend on the other attribute values. Therefore, a single-threaded multi-attribute negotiation can be equivalently transformed into multiple independent single-attribute negotiations, one for each attribute. The type of a negotiator is a vector, with one element for the reservation value on each attribute, instead of a scalar. The single-threaded negotiation model can be applied directly for negotiating a single attribute. The multi-threaded negotiations with multiple attributes are a little bit more complicated than in the situation with a single attribute, because a negotiator must trade off the offers proposed by different opponents. The offer of an opponent may not dominate the offers of other opponents on all attributes, but is better than the other offers on some

attributes. The negotiator should evaluate an offer by a scoring rule to decide whether to accept or reject an offer as a whole. An alternative to get around this difficulty is to ask negotiators to reveal their scoring rule, but leave the reservation score, the maximum/minimum score of the agreement that is acceptable, as private information. Then negotiators can directly negotiate on the score, and that transforms a multi-attribute negotiation into one single-attribute negotiation.

The same approach to negotiating a single dimensional variable can be applied to integrative multi-attribute negotiations. In an integrative multi-attribute negotiation, the scoring rule is not an additive function of the attribute values. Because of the non-linearity of the scoring function the utility associated with one attribute also depends on the values of the other attributes. This drastically increases the complexity of negotiation decisions even in a single-threaded negotiation because a negotiator has to consider how to trade-off different attributes in a proposal, and the decision changes with the negotiation status, or the latest offer and counter offer. A win-win situation, in which both negotiators get higher utility, is possible in an integrative negotiation. An outcome is *Pareto-optimal* if there does not exist an agreement that brings better utilities to both parties. The *contract curve* is the set of all agreements that are Pareto optimal. The set of all agreements that have the same utility is called a *level curve* [6]. Figure 2 shows an example of the contract curve and level curves in a integrative negotiation between negotiators a and b over attributes a_1 and a_2 .

In this example, negotiator a prefers lower values while negotiator b prefers higher values on both attributes. The intersection between the contract curve and a level curve brings the most utility to the other party among all the agreements that are indifferent to a negotiator. If an agreement is reached on the contract curve, then we can say that no efficiency is wasted. In theory, since the contract curve is the set of all Pareto-optimal agreements, the negotiators have no interest conflict in reaching an agreement on any point of the contract curve, yet they care about which agreement on the contract curve will be reached as a result of the negotiation.

An integrative multi-attribute negotiation can be transformed to a Pareto-optimal single-attribute negotiation if the proposals are restricted on the contract curve. The contract curve can be obtained by asking agents to reveal their information on how to trade-off the attributes. But again, the minimum/maximum score of an agreement that is acceptable, can be left as private information.

3.2 Risk attitude

How much a negotiator sets the reservation utility in a negotiation thread, or how much she values the outside options, depends on the risk attitude of the negotiator. For a

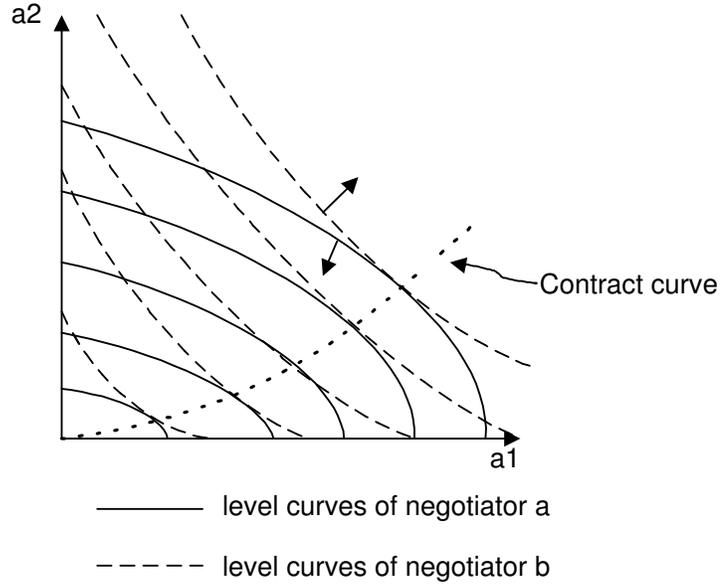


Figure 2: Contract curve and level curves

risk-neutral negotiator the reservation utility is equal to the expected utility of the outside options. But for a risk-averse negotiator the reservation utility should be less than the expected utility from outside options. This is because the actual utility from outside options is uncertain, and so the outside opportunities are devaluated by a risk-averse agent. Therefore a risk-averse agent can accept less utility from the negotiation thread under consideration than a risk-neutral agent. It is the opposite for a risk-lover.

The risk attitude of an agent can be defined by the utility transfer function $u(\cdot)$ that transfers the utility from a random event to an equivalent deterministic utility. Assume with probability $G(x)$ that the random event gives utility x . Then the agent is indifferent between the random event and a deterministic situation in which the utility is U , if and only if $U = u(\int x dG(x))$. The agent is *risk averse* if u is a concave function, i.e., $\int u(x) dG(x) < u(\int x dG(x))$, *risk neutral* if u is linear, i.e., $\int u(x) dG(x) = u(\int x dG(x))$, or a *risk lover* if u is a convex function, i.e., $\int u(x) dG(x) > u(\int x dG(x))$. We can model a wide range of risk attitude by correctly designing the utility transfer function. If a negotiator is not risk neutral, the reservation utility can be set accordingly by applying the utility transfer function to the utility distribution in the outside options, instead of using the expected value of the utility from outside options.

4 Conclusions and future work

We propose 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. Each of these three models is of increasing 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. The model of dynamic multi-threaded negotiations expands the last model by considering the uncertain outside options that may come dynamically in the future. We believe that this model provides a flexible framework to incorporate different situations, and yet reflects the essential concerns in negotiations in the Navy detailing process. The models for each modular component are described, and extensions are discussed for the application in more general situations. The modelling framework and the constituent models that have been developed lay out the foundation for our work in the next stage to analyze each modular model and to provide an integrative solution for the negotiation decision problem in the Navy detailing problem.

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] Anita Gantner. Bargaining, search and outside options. Technical report, Department of Economics, University of California, Santa Barbara, 2002.
- [5] 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.

- [6] P. Heiskanen, H. Ehtamo, and R. P. Hmlinen. Constraint proposal method for computing Pareto solutions in multi-party negotiations. *European Journal of Operational Research*, 133(1):44–61, 2001.
- [7] Pu Huang and Katia Sycara. A computational model for online agent negotiation. In *35th Annual Hawaii International Conference on System Sciences (HICSS'02)*, 2002.
- [8] 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*, 10(2):199–215, 2001.
- [9] Gregory E. Kersten and Jeffrey Teich. Are all e-commerce negotiations auctions? In *Proceedings of the Fourth International Conference on the Design of Cooperative Systems(COOP'2000)*, May 23-26 2000.
- [10] 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.
- [11] S.A. Lippman and J.J. McCall. The economics of job search: A survey. *Economic Inquiry*, XIV:155–189, 1976.
- [12] 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.
- [13] A. Muthoo. On the strategic role of outside options in bilateral bargaining. *Operations Research*, 43(2):292–297, 1995.
- [14] Roger B. Myerson and Mark A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29:265–281, 1983.
- [15] Stefan Napel. *Bilateral Bargaining: Theory and Applications*. Lecture notes in economics and mathematical systems. Springer, 2002.
- [16] Martin J. Osborne and Ariel Rubinstein. *A Course in Game Theory*. MIT Press, 1994.
- [17] Sheldon M. Ross. *Introduction to Probability Models*. Academic Press, 2000.
- [18] Katia Sycara. Utility theory in conflict resolution. *Annals of Operations Research*, 12:65–84, 1988.

- [19] Katia Sycara. Negotiation planning: An ai approach. *European Journal of Operational Research*, 46(2):216–234, May 1990.
- [20] Katia Sycara. Problem restructuring in negotiation. *Management Science*, 37(10), October 1991.
- [21] Katia Sycara. Machine learning for intelligent support of conflict resolution. *Decision Support Systems*, 10:121–136, 1993.
- [22] Daniel D. Zeng and Katia Sycara. Bayesian learning in negotiation. *International Journal of Human-Computer Studies*, 48:125–141, 1998.