A Review of Research Literature on Bilateral Negotiations

Cuihong Li, Joseph Giampapa, and Katia Sycara

CMU-RI-TR-03-41 2

A Review of Research Literature on Bilateral Negotiations

A
Carnegie Mellon University
Robotics Institute
Technical Report
CMU-RI-TR-03-41

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 11, 2003

Abstract

Automated bilateral negotiations are an important mechanism to realize efficient distributed matching in the Navy detailing system, and the presence of outside options is an outstanding feature of the negotiations. In this report we provide an extensive literature review on the research of bilateral negotiations in the fields of Economics and Artificial Intelligence. Three important dimensions are described to identify the negotiation environment and to build a research model. Preliminary considerations and suggestions are given in these dimensions on modelling the system. The review suggests that negotiations with outside options is a new and important research problem, yet it can be addressed, based on the existing work.

1 Introduction

A bilateral bargaining (negotiation) situation is characterized by two agents - individuals, firms, governments, etc. - who have a common interest in cooperation, but who have conflicting interests concerning the particular way of doing so. Bilateral bargaining refers to the corresponding attempt to resolve a bargaining situation, i.e. to determine the particular form of cooperation and the corresponding payoffs for both [32][29]. Bargaining is a prevalent form of interaction in human society. It is also an effective approach to resolve conflicts in a distributed manner, when a third party mediator is not available or trustable.

Bilateral negotiation is a useful mechanism in the Navy detailing process. In the Navy detailing process, although most matches are decided through a centralized matching market mechanism, there are situations in which commands and sailors can negotiate with each other directly and decide the matches by themselves. Please refer to Section 2 for descriptions of these situations.

The research work on bargaining has been conducted in the fields of economics, in particular game theory, and artificial intelligence (AI). The research in the economics community focuses on the outcome of a bargaining situation that satisfies some axioms[1], or the strategy equilibrium[2] 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. The complex situations considered in an AI model prohibit the strategy equilibrium to be explicitly considered. Instead the AI research pays more attention to flexible and dynamic self-optimization of an agent by learning and adapting to the environment, which also includes other agents, and searching for a good decision heuristically. Research in economics and AI have different methodologies and concerns, yet their contributions complement each other. Insights and theoretical foundations developed in economics provide good heuristics for AI, and the AI approaches provide solutions to negotiations in realistic environments that usually cannot be solved by a game-theoretic model. 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. To ground our research on bilateral negotiations, it is helpful to review the existing work and results in both the fields of economics and AI.

The rest of the report is organized as follows: Before reviewing the literature we provide background knowledge, motivate our research work, and restate our research questions in Section 2. Section 3 reviews essential vocabulary, research work and results on

---

[1] See the definition of *cooperative game theory* in Section 3

[2] See the definition of *non-cooperative game theory* in Section 3

bargaining theory in economics. The content is organized in three subjects: cooperative bargaining theory, non-cooperative bargaining theory, and bargaining with outside options. The research on bargaining in AI is reviewed in Section 4, focusing on different aspects of learning, heuristics, time issues and multi-attribute negotiations. Future work and discussions are provided in Section 5. Section 6 concludes.

2 Motivation

In the Navy detailing system a significant amount of matches are decided through a *centralized matching market.* In the centralized matching market sailors and commands submit their preference information (rankings, incentive bids, etc.), and the market replies with a matching that optimizes the overall quality across the matches. The matches are decided in a centralized way by the matching market, although sailors and commands can submit information to the market that influences the matching result. However, there are some situations in which commands and sailors can negotiate with each other and decide the matches by themselves. These situations include:

- Special jobs and sailors: Some jobs with high priorities (such as the submarine jobs) and sailors with certain specialties do not get matched through centralized matching. Matching for these sailors and jobs can be performed through bilateral negotiations between a sailor and a command.
- After centralized matching: The jobs and sailors that fail to be matched in the centralized matching market may negotiate directly with each other to reach a matching agreement.
- Direct invitation: Some sailors may not apply for a job for which they are qualified. A command may directly contact such sailors that he desires and try to make a matching agreement.

In these situations direct communication is established between a command and a sailor, and the communication may lead to an agreement between them without the mediation of a third party. We call this mechanism a *matching and bargaining market (MBM)*, which can be regarded as a substitute for the centralized matching market.

In a MBM a sailor may have multiple jobs that he qualifies for, and a command may find multiple sailors that qualify for his job. Instead of delegating the decision of matching to a centralized market, the sailors and commands negotiate one-by-one directly, and decide the matching and conditions of the matching. The matching is performed in a distributed manner since it is a result of autonomous bilateral negotiations between commands and sailors, without the coordination of a central system. Compared to the centralized matching market, the MBM enjoys high flexibility at the expense of efficiency.

The most critical problem for a bargainer during a negotiation is to decide how much to offer and how to respond (e.g., either accept or reject) to an offer. An effective negotiation decision solution is the key to automate the negotiation system supported by intelligent agents. To design an effective automated negotiation strategy, we have to pay attention to he existence of outside options, an important feature of the bilateral negotiations in the MBM. A sailor may have more than one job that he can qualify for, and a job may find multiple sailors that are qualified for it. This feature differentiates the bilateral negotiations in MBM from common bargaining problems, in which the two bargainers are monopolies. Outside options are an important factor to consider during the negotiation process because they can be used as a negotiation threat and have great impact on the negotiation strategy. With more promising outside options, a bargainer is in a more advantageous position of negotiation, and can employ a more aggressive negotiation policy. But the bargainer also takes certain risks at being aggressive, because a bargainer cannot be sure about the utility that she can achieve from the outside options. Therefore a bargainer has to tradeoff between the expectation of a future deal and the probability of losing the current chance.

The research we propose is to study the individual negotiation decisions in a matching and bargaining market by explicitly considering outside options. The task and output of this research work will be:

To build an analytical model of the negotiation problem in a MBM, and to provide a negotiation decision model and solution that maximizes the expected utility of a bargainer. The questions that will be addressed by the solution are: How should a bargainer consider the outside options in negotiations? How much should a bargainer propose to the opponent? Given an offer from the opponent, should a bargainer accept it or reject it?

3 Bilateral bargaining theory in economics

Bargaining is a type of game. Bargaining theory is a part of game theory that studies bargaining games. We use the same taxonomy of game theory, shown in Figure 1, to organize our literature review of bargaining theory.

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

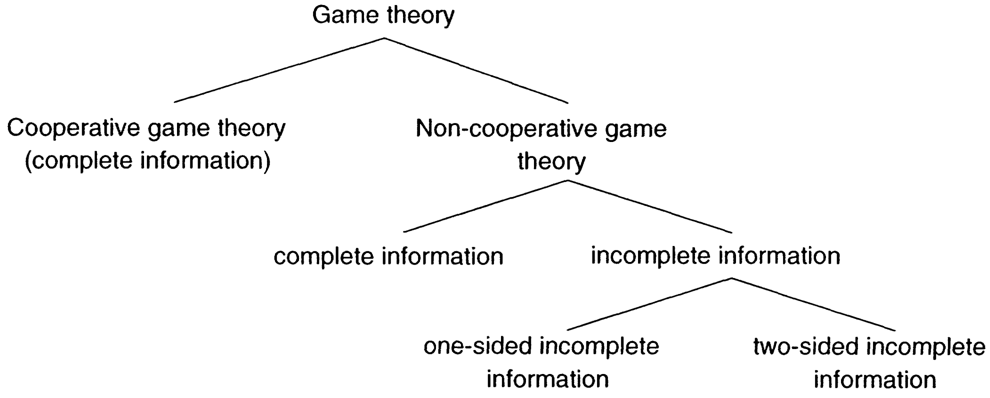

Figure 1: Taxonomy of game theory

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*. For example, an axiom can be that two bargainers get the same share of the cake that they are negotiating for. *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[3] of the negotiation. For example, the strategy of a buyer who bargains with a seller over the price can be to call the average of her last call and the seller's last call at each step, and accept an offer if it is better than the proposal she is about to submit. 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[32][34]. A strategy *equilibrium* is a profile of players' strategies so that no player could benefit by unilaterally deviating from her strategy in the profile, given that other players follow their strategies in the profile. For example, a strategy equilibrium for a buyer and seller who sequentially bargain over price without a time or budget limit is: the buyer initially calls the lowest allowable price and the seller calls the maximum allowable price. Both concede in each successive step at the minimum allowable pace, until their price calls match. No one can result in a better deal by conceding faster alone. An equilibrium facilitates the prediction of the players' behavior, and hence also the outcome of the game. Some

---

[3] 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.

widely-used concepts of a strategy equilibrium include "dominant strategy" equilibria, "Nash" equilibria and "subgame perfect" equilibria. A *dominant strategy equilibrium* consists of a *dominant strategy* of each player, which is optimal for a player irrespective of the other players' strategies. The strategies chosen by all players are said to be in *Nash equilibrium* if no player can benefit by unilaterally changing his strategy. A subgame perfect equilibrium refines the Nash equilibria in extensive-form games, i.e., games with a tree structure in which players act sequentially. In a *subgame perfect equilibrium* (SPE) the strategies for each subgame of the game tree constitute a Nash equilibrium. A game is with *complete information* if the preference information of a player is known to all other players, otherwise it is with *incomplete information.* For example, in salary negotiations between a command and a sailor, the negotiation is a *complete information* game if the sailor knows how much the command is willing to pay for the job, and the command knows the minimum salary that the sailor is willing to accept; otherwise it is an *incomplete information* game if the maximum willingness-to-pay of the command or the minimum willingness-to-accept of the sailor is private information of the command and the sailor [14][34]. If both sides have private information, it is called a *two-sided incomplete information* game, otherwise if only one side has private information, it is called a *one-sided incomplete information* game [1]. In multi-attribute negotiations the preference information also includes the relative importance of each attribute and how they trade-off with each other. For an incomplete information game, "Bayes-Nash" equilibria is the equilibrium concept that is usually used. The strategies of players, which are associated with the private information of the players, compose a *Bayes-Nash equilibrium* if no player can get higher benefit *on expectation* by unilaterally changing his strategy [14][34]. Cooperative games are all based on complete information, assuming that the input to the axiomatic solution is common knowledge or that players share true information with each other and the mediator, who regulates the solution given the information. In non-cooperative game theory players may withhold information or not be truthful with each other [34].

3.1 Cooperative bargaining theory

Cooperative bargaining theory is concerned with the question of what binding agreement two bargainers would reach in an unspecified negotiation process given the set of all possible agreements on the utility that each bargainer achieves. The path-breaking work of Nash [33] provides a unique solution that satisfies four properties, which are now called the "Nash axioms". The Nash axioms include: (1) The final outcome should not depend on how the players' utility scales are calibrated; (2) The agreed payoff pair should always be *individual rational*[4] and *Pareto-efficient*[5]; (3) The outcome should be independent of irrelevant alternatives; (4) In symmetric situations, both players get the same utility [16]. 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. Some other solution concepts in cooperative bargaining theory include the *Kalai-Smorodinsky bargaining solution* [20] and *weighted utilitarian (bargaining) solution* [30][32]. A cooperative bargaining model does not consider the negotiation process, but leaves the outcome to be determined by an axiom. We do not expect cooperative bargaining theory to be applicable to our research on bilateral negotiations in Navy detailing systems because the negotiation outcome is the result of specific interactions between commands and sailors.

3.2 Non-cooperative bargaining theory

Non-cooperative bargaining theory considers bargaining as a fully specified game. The game refers to the negotiation protocol that two players follow during the bargaining process. A negotiation protocol is the set of rules that govern the interaction between negotiators. It covers the negotiation states (e.g. accepting proposals, negotiation closed), the events that cause negotiation states to change (e.g. no more bidders, bid accepted), and the valid actions of the participants in particular states (e.g. which messages can be sent by whom, to whom, at what stage) [19]. Some negotiation games or protocols are described below. These protocols have been applied mainly to evaluate negotiations over a single issue, such as the price of a good to be negotiated. They can be extended to the multi-attribute negotiations in which the attributes are negotiated either simultaneously or sequentially. In the following we first describe the negotiation protocols and present the subgame perfect equilibium (SPE) of the complete information bargaining game under these protocols. Then in the rest and most of the section we review the work in non-cooperative game theory with incomplete information.

3.2.1 Bargaining with complete information

We describe three non-cooperative bargaining games with complete information. These games can be used by two bargainers to divide a given bargaining "surplus" - the total profit resulting when the players reach an agreement. The total surplus resulting from

---

[4] An agreement is individual rational if a player does not lose by participating in the game.

[5] An agreement is Pareto-efficient if no player can gain without causing a loss for the other player.

the match between a sailor and a job can be equated to the value of the sailor to the job, minus the minimum benefit required by the sailor. These games are different in the negotiation protocols, hence also the outcome at the subgame perfect equilibria. Such protocols can be candidates of the bargaining protocols used in the Navy detailing negotiation system. The subgame perfect equilibria of these games provide valuable references to design the autonomous negotiation strategies if corresponding protocols are used, and complete information is assumed.

The ultimatum game

One of the players proposes a split of the surplus and the other player has only two options: accept or refuse. In case of refusal, both players get nothing. In this game the proposer has overwhelming bargaining power. The subgame perfect equilibrium (SPE) specifies that the proposer asks for the entire surplus, and the responder accepts [32].

The alternating-offers game

The alternating-offers game is a multi-stage extension of the ultimatum game. In this game player 1 starts by offering a fraction x of the surplus to player 2. If player 2 accepts player 1's offer, he receives x and player 1 receives $1 - x$ of the surplus. Otherwise, the game moves to the second stage, and player 2 needs to make a counter offer. Player 1 can accept, or reject and make another proposal in the next stage. This process is repeated until one of the players agrees or until a finite deadline is reached. A unique SPE can be found with backward induction. The SPE reaches an immediate agreement on an efficient division of the surplus. The driving force behind this result is players' impatience to reach an agreement and one player's opportunity to make a take-it-or-leave-it offer in the final period [32].

Rubinstein [39] studies the alternating-offers game with infinite horizon, i.e., the bargaining can go on forever if no player accepts an offer of the other player. This model with infinite horizon can be used to approximate the situation with many negotiation rounds available. It shows that the strong result of a unique SPE prediction can be generalized from the finite case to that of infinitely many stages. The unique SPE approaches the outcome of the Nash bargaining solution, as described in Section 3.1, with the bargaining power of the two players $\alpha = \frac{\ln\delta_2}{\ln\delta_1+\ln\delta_2}$, $\beta = 1 - \alpha$, where δ_1 and δ_2 are the time discount factors of player 1 and player 2 [32].

The monotonic concession game

In the monotonic concession protocol, there are two players, and they announce their proposals simultaneously. If the offers of both players match or exceed each other's demands, an agreement is reached, and a coin is tossed to choose one of the offers. If no agreement is reached, the players need to make new offers in the next round. The offers need to be monotonic, i.e., the new offer should not have a lower utility for the counter player compared to the last offer. Hence, a player can either make the same offer or concede. The negotiation ends if both players stand firm in the same round, and players leave without an agreement [16]. For details on the analysis of this game, please refer to [38].

3.2.2 Bargaining with incomplete information

The models introduced in Section 3.1 and Section 3.2.1 are all based on complete information, i.e., the players know each other's preferences. More commonly and realistically a bargainer does not know exactly how the counter bargainer evaluates the possible agreements while they negotiate over the surplus distribution or price. Some other important information, such as the attitude towards risk, time preferences, and preferences among issues, are also often hidden from the opponent in real-life negotiations. This situation is referred to as *bargaining with incomplete information.* The private information that is unknown to the the other players in the game is usually referred to as the *type* of the player.

For convenience of explanation and ease of understanding, we can call the bargainers a buyer and a seller. The seller negotiates with a buyer as to the price of an indivisible good. The seller values the good at c, which can be the cost of producing the good. The buyer associate some value v to the good.

In the complete information negotiation models, the equilibrium strategies imply immediate agreement because the benefit is discounted with the time that is elapsed in the negotiation process. But in negotiations with incomplete information, an agreement may be delayed as the parties continue negotiations, or might never be reached. Delays are costly but may be required to convey private information credibly, and to screen or signal the types of the players (Kennan & Wilson [22]).

Research on bargaining with incomplete information can be conducted in a *mechanism design* framework or in a *sequential bargaining* game framework. Ausubel, et al. [1] provide a survey from both perspectives. In a *mechanism design* framework the process of bargaining is abstracted away. Rather than model bargaining as a sequence of offers and counteroffers, the bargaining mechanism is analyzed as a mapping from the parties' private information to bargaining outcomes. This permits the identification of properties shared by all Bayes-Nash equilibria of any bargaining game. Myerson and Satterthwaite (1983) [31] show that whenever the two players are uncertain about each other's type, there are always some pairs of the player types that do not reach an agreement although an agreement is actually desirable for both parties. For example, if the seller does not know the buyer's valuation v and the buyer does not know the seller's cost c, then there are some v and c such that the seller and the buyer cannot reach a deal, although $v > c$. The inefficiency comes with the incomplete information. Since players are not willing to show their reservation prices but expect the concessions of the other, there are some situations where they leave without an agreement although they should have been able to reach an agreement if they were truthful.

Mechanism design provides a powerful theory for studying incentive problems in bargaining, characterizes the set of outcomes that are attainable, and determines optimal trading mechanisms. But the mechanisms studied in mechanism design are static and mediated mechanisms, in which players reveal their types and then an outcome is generated. These mechanisms require the commitment that bargainers in some situations have to walk away without an agreement, although they both know that an agreement would have benefited both after they reveal their types.

Sequential bargaining games are more interesting in practice, and we believe, to the Navy detailing system, because the trading result is decided by a process of offers and responses of the traders [1]. The work on sequential bargaining games with incomplete information can be based on *one-sided* or *two-sided* incomplete (private) information. In *one-sided private information* bargaining, only c or v, but not both, is private information, whereas in *two-sided private information* bargaining, both c and v are private information. In the following subsections we describe the main research results on sequential bargaining with one-sided or two-sided private information. The major conclusive work include Fudenberg and Tirole (1983) [13] for a two-period model with one-sided and two-sided incomplete information, Rubinstein (1985) [40] for one-sided incomplete information with an infinite time horizon, and Chatterjee and Samuelson [5][6] for two-sided incomplete information in an infinite time horizon model with two types.

One-sided incomplete information

Without loss of generality we assume that the informed party is the buyer, i.e., the buyer knows c and v, but the seller only knows c. Let the seller's cost c be zero, for the ease of presentation. The three major protocols considered are: the seller-offer game, the alternating-offers game, and the buyer-offer game. In a *seller-offer* or *buyer-offer* game only the seller or buyer is the one to propose offers, and the other party can only respond by accepting or rejecting an offer. In an *alternating-offer* game the two parties alternate their roles of proposing and responding to offers.

In a **seller-offer game**, the uninformed party, the seller, proposes offers and the informed party, the buyer, accepts or rejects. In this situation the buyer's acceptance behavior is completely characterized by a non-increasing acceptance function. Consequently, the seller's belief on the buyer's type will always be a truncation of the prior belief [1]. This means that the set of the buyer's types that the seller thinks possible, will shrink. The Bayes-Nash equilibrium can be calculated by dynamic programming. The solution of the game follows Sutton [44]. Suppose the game terminates after T periods. For each T there exists a unique sequential equilibrium. If the seller's prior belief about v is described by a uniform distribution on $[0, 1]$, then as $T \to \infty$, the sequential equilibrium converges to an equilibrium sequence of the infinite horizon game. In the resulting equilibrium the seller calls a price that is a constant fraction of the previous call.

In a seller-offer game, the informed party has very limited means of communication, and can only be screened by two possible responses: accept the current offer and thereby terminate the game, or reject the current offer in order to trade at more favorable terms in the future. In an **alternating-offer game**, the informed party can also propose, and so has a much richer language with which to communicate [1]. There is a potential in the alternating-offer game that the buyer signals her type, with higher valuation buyer types trading-off higher prices for a higher probability of acceptance. Ausubel and Deneckere (1998) [3] propose a refinement of perfect equilibrium, termed *assuredly perfect equilibrium* (APE). Intuitively APE requires that beliefs should be concentrated on the lower valuation buyer types if there is an unexpected move by the seller, or the lower valuation buyer types are more likely to deviate from the equilibrium strategy than the higher valuation buyer types. Based on APE they show that the alternating-offer game requires more offers, and has a lower acceptance probability than the seller-offer game. Moreover, the alternating-offer game results in additional delay.

In a **buyer-offer game** there is only one sequential equilibrium, which is for the buyer to always offer the seller the cost c, and for the seller to accept any price above c with probability one. It is because the buyer knows the minimum price that the seller can accept, which is equal to the seller's cost c. The buyer then will never call a price higher than c.

Two-sided incomplete information

The theory of mechanism design shows that it is not possible to achieve *ex post* efficiency[6] when private information is present. An interesting question to ask is: Can

---

[6] A mechanism is *ex post efficient* if the outcome maximizes the total utility of players given players' types. For example, a bargaining mechanism between a buyer and a seller is ex post efficient if and

the optimal static bargaining mechanism that maximizes the total utility of the bargainers on expectation be replicated in a sequential offer/counteroffer bargaining game? Ausubel and Deneckere (1993) [2] establish that, for distribution functions exhibiting monotonic hazard rates[7] the optimal static bargaining mechanism *can* essentially be replicated in very simple sequential bargaining games that use a seller-offer or buyer-offer protocol. This conclusion states that there need not be any additional inefficiency arising from the dynamic nature of the game beyond the inefficiency already introduced by the two-sided incomplete information. Chatterjee and Samuelson [5][6] analyze the equilibrium strategy of a bargaining game with two-sided private information, assuming each party has only two types.

3.3 Bargaining with outside options

Usually the result of bargaining depends on the circumstances, which often include outside options. Outside options are the alternatives that a bargainer may take if the current bargaining fails. On the trading markets, the outside options of a buyer can be the price quotes that are offered by other sellers. The presence of outside options is an important feature of the Navy detailing market. When a command negotiates with a sailor, the negotiation strategies and outcome are influenced by the possible existence of other candidates with different backgrounds. Similarly, a sailor considers the potential offers that come from other commands.

Binmore, et al. [4] establish the Outside Option Principle, which states that only those outside options with payoffs that are superior to Rubinstein's equilibrium payoffs have effect on equilibrium strategies, and this is important because in some cases it yields an equilibrium payoff different from a split-the-difference outcome. Rubinstein and Wolinsky (1985) [41] incorporate an outside option in an alternating-offer bargaining process. Chatterjee and Lee (1993) [7] study the situation where there is incomplete information about the outside option. Cunyat [8] re-examined the robustness of the Outside Option Principle based on Rubinstein's bargaining model with complete information. His paper argued 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 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 [28] studies a model of the situation in which two players are bargaining face-to-face over the partition of a cake,

---

only if it results in a trade whenever the buyers' valuation is greater than the seller's cost.

[7] With monotonic hazard rates, $\frac{1-G(v)}{g(v)}$ is a strictly decreasing function, $\frac{F(c)}{f(c)}$ is a strictly increasing function.

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[8].

All the above cited papers in this section are based on complete information about the bargainers' preference. Gantner [15] extends the model with incomplete-information and alternating-offers in Chatterjee and Samuelson (1987) [5] by introducing an outside option which is modelled as a standard sequential search process, where the buyer can choose to search for other offers and return to bargaining at any time. The outside option is described by an arrival of offers according to a Poisson process. The searching policy and bargaining strategies are considered jointly.

4 Bilateral bargaining research in AI

The goal of game theoretic research on bargaining is to identify the optimal mechanisms with desirable properties, and to analyze the strategy equilibrium with a bargaining protocol. Game theory usually assumes complete knowledge of the circumstances, and full rationality of the players. The first assumption implies that the rules of the games and beliefs of the players are "common knowledge"[9]. The second assumption implies that players have infinite reasoning and computational capacity to maximize their expected payoffs given their beliefs of others' types, behaviors and beliefs[10]. These two assumptions limit the practical applicability of game-theoretic results. In the field of AI, however, these two assumptions are not necessary and agents are not expected to play exactly following equilibrium strategies. The emphasis lies more on finding acceptable rather than optimal solutions. The AI approaches are able to handle more complex situations but at the same time with more simplified models in a more practical way than the more rigorous game-theoretic models. However, the connection between the AI approach and game-theoretic analysis is important. Game theory can provide valuable managerial insights, and aid in the difficult task of choosing a suitable bargaining protocol. AI techniques can be used to develop software applications and bargaining protocols that are currently beyond the reach of classic game theory [16]. Empirical and theoretical study has found that learning by repeatedly playing games also leads to an equilibrium [21].

---

[8] See the definition of a "Nash bargaining solution" in Section 3.1.

[9] Even with incomplete information, the prior distribution of the players' types are common knowledge.

[10] The beliefs are statements like "I know that he knows that I know ... " ad infinitum.

4.1 Learning in bargaining

In AI, learning is an important technique to deal with an environment of which complete knowledge is not known. Knowledge acquired by learning can further enable an agent to behave flexibly and to adapt to changes in their environment. In bargaining games, learning is potentially important in the following two aspects [16]. First, a bargaining agent can adjust its bargaining strategy so as to achieve better deals, based on its experiences in previous bargaining games. Second, it is useful to update the belief of the other parties' types or strategies during play by observing the behaviors of the other parties, and then to adjust one's strategy accordingly. Gerding [16] introduces several AI learning techniques including decision trees, Q-learning, evolutionary algorithms and Bayesian beliefs. Emphasis is given to applying evolutionary approaches and Bayesian beliefs to learn effective strategies or useful information in negotiations. Zeng and Sycara [46] present a sequential negotiation model and address multi-agent learning issues by explicitly modelling beliefs about the negotiation environment and the participating agents under a probabilistic framework using a Bayesian learning representation and updating mechanisms. In Mudgal and Vassileva [27] agents use influence diagrams to create models of their opponents during negotiation, which help them to better predict their opponents' actions.

4.2 Negotiation heuristics

Usually searching for a strategy equilibrium is computationally untractable, unless an analytical solution is known. This is because an equilibrium involves outguessing regress. *Outguessing regress* means that the decision of an agent depends on its beliefs about other agents, each of which in turn needs beliefs about the other (former) players, and so on. In reality agents have bounded rationality, that is, their computation and reasoning resource is limited. Therefore they are usually not expected to play a game following the equilibrium strategy. AI approaches help the players locate an approximate solution strategy according to principles of bounded rationality by utilizing heuristic search and evaluation techniques [43][19].

Faratin, et al. [10] define a range of computationally tractable heuristic strategies and tactics that negotiating agents can employ to generate initial offers, evaluate proposals and offer counter proposals in multi-attribute negotiations. The *tactics* are simple functions that are used to generate an offer, or counter offer, based on different criteria. A *strategy* is the way in which an agent changes the weights of the different tactics over time. In the negotiation model agents propose offers alternatively following their strategies. Each agent has a scoring function that is used to rate the offers received. If an agent receives an offer that has value greater than the value of the counter offer that it is ready to submit in the next step, then it accepts. Otherwise, the counter offer is submitted. The negotiation tactics include time-dependent tactics, resource-dependent tactics and behavior-dependent tactics.

In the **time-dependent tactics**, an agent submits offers that change monotonically from the minimum (best) to the maximum (worst) of the deal that she can agree on, and the rate of change depends on time. Faratin, et al. [10] distinguish two families of the changing rate functions, with the rate of change being a polynomial or an exponential function of time. By varying the parameters, the functions can model that agent to be more of a boulware [37, pg. 48] or a conceder [36, pg. 20]. A bargainer is a *boulware* if she does not concede until close to the last moment, and a bargainer is a *conceder* if she gives up quickly.

The **resource-dependent tactics** are similar to the time-dependent ones in which time is the sole considered resource. The resource-dependent tactics are modelled in the same way as the time-dependent ones by using the same functions. The difference is that the resource-dependent tactics either, (i) have dynamic value of the maximum available resource, or (ii) make the changing rate function depend on an estimation of the amount of a particular resource.

The **behavior-dependent tactics** compute the next offer based on the previous attitude of the negotiation opponent. These tactics are especially important in cooperative problem solving negotiation settings, or integrative negotiations, by allowing agents to consider the other agents' behavior. The main difference between the tactics in this family is in the type of imitation they perform. One family imitates proportionally, another in absolute terms, and the last one computes the average of the proportions in a number of previous offers.

Huang and Sycara [18] present a formal autonomous negotiation model. In this model, the negotiation process is driven by the internal beliefs of participating agents following a group of belief updating methods. The "personality" of an agent can be realized by using appropriate belief updating methods with suitable parameters. If an agent is "tough" then she will believe more as time progresses that the other agent will concede. If an agent is "weak" then she becomes more convinced as time progresses that the other agent will not change its current offer . In this paper the belief of the other agent's type takes a uniform or exponential probability distribution. With different belief updating methods, the probability distributions have different supporting space, and different skewing directions and slopes for exponential distributions.

Note that although the heuristics in the two papers mentioned above provide flexible and formal ways to compose negotiation strategies, the strategic interactions between agents are not explicitly considered, i.e., agents do not speculate on the possible responses of the other agents when they decide their actions according to these heuristics.

4.3 Time issues in bargaining

Some work in AI has considered the time issues in negotiations. Generally speaking the most common effects of time on the bargaining process are [11]:

- Discounting: Benefits received immediately are preferred to the same benefits received in the future
- Bargaining cost: The bargaining process itself incurs some cost/utility to an agent
- Sudden termination: There is a hard deadline beyond which the negotiation cannot be continued or is useless

Kraus, et al. [25] propose a strategic model of negotiation that takes into account the passage of time during the negotiation process, itself. Some agents may lose over time while others gain over time, such as in a negotiation for resource sharing between an agent that currently occupies the resource and an agent that is in need of that resource. The loss and gain of value or cost as a result of time is private information of an agent. Further, agents are not required to negotiate until a deadline, and may leave a negotiation at will. The equilibrium strategies of the agents are based on game-theoretic analysis, and result in efficient agreements without delays. Sandholm and Vulkan [42] analyze automated distributed negotiation where agents have firm deadlines that are private information. It shows that the only sequential equilibrium outcome is one where the agents wait until the first deadline, at which point the agent with the earlier deadline concedes everything to the other. Fatima, et al. [11] analyze the optimal negotiation strategies for autonomous bargaining agents that have firm deadlines and incomplete information about their opponent. The model sheds insights on how an agent can exploit its available partial information to make decisions in the optimal strategy.

4.4 Multi-attribute negotiations

Sometimes multiple attributes are involved in negotiations, such as the multiple terms of a supply contract or the different properties of a product [19][23]. In a negotiation with a single attribute, the gain of one player always creates loss of the other player.

This kind of negotiation is referred to as "competitive". When multiple criteria are involved, and players attach different importance to each criterion, there may exist some outcomes that are better for both parties than the current offer. This is called an "integrative" negotiation. We should note that not all multi-attribute negotiations are integrative. If the utility function of an agent is the weighted sum of all attributes, the negotiation can be equivalently transformed into multiple independent single-attribute negotiations, one for each attribute. This is possible because the utility associated to one attribute is independent of the values of the other attributes when the preference function is in a linear form.

Unfortunately, with the exception of Ponsati and Watson [35], and Fershtman [12], there has been little work on game theory that studies multi-attribute negotiations. Fershtman [12] considers sequential bargaining for two criteria and analyzes the optimal agenda for a negotiator. Ponsati and Watson [35] present axiomatic solutions for a multi-attribute negotiation instead of solutions based on sequential bargaining protocols. In an axiomatic solution the outcome of the negotiation is determined by a mapping from the preferences of the negotiators to the agreements that satisfy certain properties, without specific negotiation processes.

Some AI heuristics have been proposed to achieve win-win situations in multi-attribute negotiations. An important consideration in the heuristics is how an agent can use the preference information that is revealed by the other agent in its proposal generation, in such a way that the improvement of one agent's interest is not necessarily in the cost of sacrificing the other's. Faratin [9] proposes the use of similarity criteria to make negotiation trade-offs. The idea is for an agent to make a counter-offer that maximizes the similarity to the opponent's last offer among all the agreements that are on the same certain utility level. Heiskanen, et al. [17] present a constraint proposal method for computing Pareto optimal solutions in multi-party negotiations. A solution is *Pareto optimal* if there does not exist an agreement that brings better utilities to all parties. In this model, agents are cooperative and propose offers alternatively, following the direction from the current offer that was suggested by the other party. In both [9] and [17] the information about the preference trade-off of the other party is used in an agent's decision to make an offer in a cooperative way. In [9] the information is revealed by the last offer made by the other party. In [17] the information is abstracted in the derivative of the other party's utility function at the current offer.

Rather than propose cooperative search heuristics as in [9] and [17], Klein, et al. [24] describe a simulated annealing approach to negotiate complex contracts with a mediator. In Tesauro [45] and Li and Tesauro [26] a formative strategic decision model for bilateral multi-attribute negotiations under incomplete information is presented for the one-side offer and alternating-offer protocol. The strategy computes offers (and offer acceptance when an alternating-offer protocol is used) to approximately optimize expected utility using depth-limited combinatorial search and Bayesian updating. The optimization is performed with respect to a model of the opponents' presumed strategic behavior and a probability distribution on the opponent's utility function. The probability distribution is updated based on observed offers and responses. The advantage of this model is that it combines learning and optimization, and allows strategic decisions by specifically speculating on the other agent's responses and self gains in the future. The disadvantage is that the speculation on the other agent is based on the presumed strategic behavior of the other agent, and so lacks flexibility. The computation can be prohibitive when the search and reasoning depths are big.

5 Future work and discussions

In this report we have reviewed the existing work on bilateral negotiations. Our future work includes modelling and analyzing the negotiation problem in the Navy detailing system. To make the transitions possible to the next task, we need to identify the situations in the following three important dimensions:

- Information: Will the negotiations between commands and sailors be with complete information, one-sided incomplete information, or two-sided incomplete information? In other words, to what extent will each side be aware of the other side's: preferences, the best agreement that they can offer, or what they are likely to accept?

- Protocol: The protocol decides the roles and actions that the negotiator can take at each moment of the negotiation process. First, it should be decided if the negotiation is a one-shot or a sequential process. Second if the negotiation is a sequential process, is it an ultimatum protocol, in which only the sailor or the command can propose offers and the other party can only respond by accepting or rejecting the offers, or is it an alternating-offer protocol, in which the sailor and the command propose and respond in turns?

- Threads: A thread is a negotiation between a sailor and a command for one job assignment. Because of the existence of outside options, a sailor can negotiate for more than one job, or a command for more than one sailor. Such negotiations are called *multi-threaded* if a negotiator can be involved in more than one negotiation, simultaneously, for the same goal, but only if an agreement can only be reached in at most one of the negotiations. Otherwise, if a negotiator can participate in only one negotiation for a goal at one time, then it is called a *single threaded* situation. Note that the multiple threads have to be associated with the same goal. For example, if a command is involved in two negotiations at the same time, but for different positions, it is still called a single-threaded situation. A thread is *resumable* if a negotiator can leave a negotiation for a while, and then return and resume it. This can be the situation where a negotiator suspends the negotiation to search for a better deal outside. A resumable thread allows the negotiator to return and continue the negotiation if a better deal is not located. In the *single* and *un-resumable* threaded situation, a negotiator must finish one negotiation before she starts a negotiation with another alternative partner. In the *single* and *resumable* threaded situation a negotiator can switch between multiple negotiations, but there can be only one thread active at one time. Similarly, in the *multiple* and *resumable* threaded situation, a negotiator can temporarily suspend some threads while pushing forward with the others. But in a *multiple* and *un-resumable* threaded situation, all threads must be active.

Based on our understanding and analysis of the real situation, we make the following preliminary suggestions to limit the space of considerations:

- Information: The negotiations are in the two-sided incomplete information situation. Usually a command does not know a sailor's personal preferences for incentive features. Although a sailor may have information about the limits to which a command can offer standard incentive features based on Navy regulations, the real limit may depend on the specific situation and budget of individual commands, which are unknown to the sailors. A command may also offer incentives on some non-standard or flexible features for which there are no regulations.

- Protocol: Sequential negotiations are preferred to one-shot negotiations. One-shot negotiations limit the negotiators' chances of adjusting the offers, and also require complicated reasoning and computation. Sequential negotiations, however, allow negotiators more decision flexibility and reduce the decision complexity by providing more information about the parties during the negotiation procedure.

  The sequential negotiation can be in the form of a command-offer protocol, or alternating-offer protocol. A command-offer protocol is plausible because a sailor may not have as much negotiation power as a command, especially when the sailor is at a relatively junior level of seniority. It is also reasonable if the sailor is the party with more information than the command. The alternating-offer protocol is worth consideration, however, if the two parties are somewhat symmetric in their bargaining power and information.

- Threads: A negotiator can participate in multiple threads, but we presume that the maximum number of threads is small in number. Typically, a bilateral negotiation is initiated by a command attempting to recruit a sailor to a position. Since a sailor does not initiate a negotiation, the chance that a sailor is involved in many threads at the same time is small. The number of threads for a command is also small because a command only initiates a bilateral negotiation if the position is not, or cannot be filled through the mass-matching market. In these situations the qualified candidates are few, and are usually discovered sequentially and dynamically by search effort.

  We would like for the threads to be non-resumable in order to reduce the complexity of our initial analysis. Prior research on bargaining [15, 28] shows that having the option to suspend and resume a negotiation does not improve the utility of a negotiator in typical bargaining situations. Also, the assumption that there will be outside options available to either party reduces the likelihood that the waiting party will remain committed to a suspended thread.

In order to propose an effective negotiation strategy solution we must first build a model to characterize the negotiation situation, which is not possible without first identifying the situations in the above three important dimensions. Irrespective of the uncertainty of the model at this point, we can see, based on the review and the preliminary considerations on the model, that no existing work can answer the question, *what is the optimal or effective strategy of a bargainer when she faces uncertain outside options.* Although there has been some work in game theory that studies bargaining with outside options, most of it is based on complete information. The one paper that studies bargaining based on incomplete information (Gantner [15]) assumes that agents have only two types, and even with such a simplified model, still fails to give an idea of the optimal strategy. We have not seen any research work in AI that studies bargaining with outside options explicitly considered, although there is potential to adapt the negotiation heuristics and techniques that have been developed in the prior work. Bilateral negotiations with outside options is a new and important research problem from both theoretical research and practical perspectives. We believe that the existing work in economics provides valuable references on modelling and analyzing the impact of outside options on negotiations, and that the work in AI will help in the design of effective heuristic negotiation strategies.

Note that some of the work reviewed in this report, such as cooperative bargaining game theory, bargaining theory with complete information, etc., may not be related or not useful to our immediate research on bilateral negotiations in the Navy detailing process. This work is included in this report for completeness, and more importantly, for considering their potential value as references in possible future work with different situations or extensions.

6 Conclusion

Bilateral negotiations are an important mechanism to realize distributed matching in the Navy detailing system. In this report we define the basic concepts and provide a vocabulary in bargaining theory for communications of research and ideas. An extensive research literature review is given on bilateral negotiations in both fields of economics and AI, which complementarily contribute to the bargaining research. We describe three dimensions that are important to identify the situations and model the system, and provide preliminary considerations along these dimensions. Existing work suggests that bilateral negotiations with outside options is a new and valuable research problem.

Acknowledgements

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

References

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

[2] Lawrence Ausubel and Raymond Deneckere. Efficient sequential bargaining. *Review of Economic Studies*, 60:435–462, 1993.

[3] Lawrence Ausubel and Raymond Deneckere. Bargaining and forward induction. Technical report, University of Maryland and University of Wisconsin-Madison, 1998.

[4] Ken Binmore, Avner Shaked, and John Sutton. An outside option experiment. *Quarterly Journal of Economics*, 104(4):753–770, 1989.

[5] K. Chatterjee and L. Samuelson. Bargaining with two-sided incomplete information: an infinite horizon model with alternating offers. *Review of Economic Studies*, LIV:175–192, 1987.

[6] K. Chatterjee and L. Samuelson. Bargaining under two-sided incomplete information: the unrestricted offers case. *Operations Research*, 36:605–618, 1988.

[7] Kalyan Chatterjee and Ching Chyi Lee. Bargaining and search with incomplete information about outside options. *Games and Economic Behavior*, 22:203–237, 1998.

[8] Antoni Cunyat. Strategic negotiations and outside options. Technical report, Department of Economic Analysis, University of Valencia, 1998.

[9] P. Faratin. Using similarity criteria to make negotiation trade-offs. In *Proceeding of the 4th International Conference on Multi-Agent Systems*, pages 119–126, Boston, USA, 2000.

[10] 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.

[11] S. S. Fatima, M. Wooldridge, and N. R. Jennings. Optimal negotiation strategies for agents with incomplete information. In John-Jules Meyer and Milind Tambe, editors, *Pre-proceedings of the Eighth International Workshop on Agent Theories*, pages 53–68, 2001.

[12] C. Fershman. The importance of the agenda in bargaining. *Games and Economic Behavior*, 2(224-238), 1990.

[13] D. Fudenberg and J. Tirole. Sequential bargaining under incomplete information. *Review of Economic Studies*, 50:221–248, 1983.

[14] Drew Fudenberg and Jean Tirole. *Game Theory*. MIT Press, 1991.

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

[16] E. H. Gerding, D.D.B. van Bragt, and J.A. La Poutré. Scientific approaches and techniques for negotiation: a game theoretic and artificial intelligence perspective. Technical Report SEN-R0005, CWI, 2000.

[17] 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.

[18] Pu Huang and Katia Sycara. A computational model for online agent negotiation. In *35th Annual Hawaii International Conference on System Sciences (HICSS'02)*, 2002.

[19] 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.

[20] E. Kalai and M. Smorodinsky. Other solutions to Nash's bargaining problem. *Econometrica*, 43:513–518, 1975.

[21] Ehud Kalai and Ehud Lehrer. Rational learning leads to Nash equilibrium. *Econometrica*, 61(5):1019–1045, 1993.

[22] John Kennan and Robert Wilson. Bargaining with private information. *Journal of Economic Literature*, 31(1):45–104, 1993.

[23] 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.

[24] Mark Klein, Peyman Faratin, Hiroki Sayama, and Yaneer Bar-Yam. Negotiating complex contracts. *Group Decision and Negotiation*, 12(2), 2003.

[25] Sarit Kraus, Jonathan Wilkenfeld, and Gilad Zlotkin. Multiagent negotiation under time constraints. *Artificial Intelligence*, 75(2):297–345, 1995.

[26] Cuihong Li and Gerald Tesauro. A strategic decision model for multi-attribute bilateral negotiation with alternating offers. In *Proceedings of the 4th ACM Conference on Electronic Commerce*, pages 208–209, June 2003.

[27] C. Mudgal and J. Vassileva. Bilateral negotiation with incomplete and uncertain information. In *Proceedings of the Workshop on Cooperative Information Agents*, 2000.

[28] A. Muthoo. On the strategic role of outside options in bilateral bargaining. *Operations Research*, 43(2):292–297, 1995.

[29] A. Muthoo. *Bargaining theory with applications*. Cambridge University Press, 1999.

[30] R.B. Myerson. Utilitarianism, egalitarianism, and the timing effect in social choice problems. *Econometrica*, 49(4):883–897, 1981.

[31] Roger B. Myerson and Mark A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29:265–281, 1983.

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

[33] J. F. Nash. The bargaining problem. *Econometrica*, 28:155–162, 1950.

[34] Martin J. Osborne and Ariel Rubinstein. *A Course in Game Theory.* MIT Press, 1994.

[35] Clara Ponsatí and József Sákovics. Rubinstein bargaining with two-sided outside options. *Economic Theory*, 11:667–672, 1998.

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

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

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

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

[40] A. Rubinstein. A bargaining model under incomplete information. *Econometrica*, 53:1151–1172, 1985.

[41] A. Rubinstein and A. Wolinsky. Equilibrium in a market with sequential bargaining. *Econometrica*, 53:1133–1150, 1985.

[42] T. Sandholm and N. Vulkan. Bargaining with deadlines. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 1999.

[43] H. A. Simon. *Models of Bounded Rationality*, volume 2. MIT Press, 1982.

[44] John Sutton. Non-cooperative bargaining theory: an introduction. *Review of Economic Studies*, 53(5):709–724, 1986.

[45] G. Tesauro. Efficient search techniques for multi-attribute bilateral negotiation strategies. In *Proceedings of the Third International Symposium on Electronic Commerce (ISEC-02)*, pages 30–36. IEEE Computer Society, 2002.

[46] Daniel D. Zeng and Katia Sycara. Bayesian learning in negotiation. *International Journal of Human-Computer Studies*, 48:125–141, 1998.