Decision Making of Negotiation Agents Using Markov Chains

Bo An*, Kwang Mong Sim†, Chun Yan Miao‡ and Zhi Qi Shen§

August 12, 2006

Abstract

In dynamic and complex negotiation environments, a negotiation agent can participate or quit negotiation at any time and can reach an agreement with more than one trading partner as the result of the existence of dynamic outside options. Thus, it's important for a negotiation agent to make a decision on when to complete negotiation given market dynamics. Rather than explicitly modelling all the trading partners, this paper presents a novel decision making strategy that an agent can use to determine when to complete negotiation based on a tractable Markov chain model of negotiation process. Experimental results suggest that the proposed strategy achieved more favorable negotiation outcomes as compared with the general strategy.

1 Introduction

Automated negotiation [12, 14] among software agents is becoming increasingly important as it enables (self-interested) agents to find agreements and partition resources efficiently and effectively. Widely studied in Economics, Law and Artificial Intelligence, research in automated negotiation is receiving increasing attention. One key challenge of automated negotiation is to design negotiation agents that can achieve good performance in dynamic, complex environments closed to real world domains. Even though there are many extant negotiation agents for e-commerce (e.g., [3, 5]), the strategies of some of these agents are mostly static and may not necessarily be the most appropriate for changing market situations since 1) they seldom consider the dynamics of the market and/or 2) they assume that an agent often has complete information about others.

This work focuses on automated negotiation in dynamic and complex negotiation environments which often have the following three characteristics: 1) an agent can

*Bo An is with the Computer Science Department, University of Massachusetts, Amherst, MA 01003 USA (e-mail: boancqu@gmail.com).

†Kwang Mong Sim is with the Department of Computer Science, Hong Kong Baptist University, Kowloon Tong, Hong Kong (e-mail: bsim@comp.hkbu.edu.hk).

‡Chunyan Miao is with School of Computer Engineering, Nanyang Technological University, Singapore (e-mail: ascymiao@ntu.edu.sg).

§Zhiqi Shen is with School of Electrical & Electronic Engineering, Nanyang Technological University, Singapore (e-mail: zqshen@ntu.edu.sg).

reach an agreement with more than one trading partner, i.e., there are outside options in negotiation. Outside options can increase agents' bargaining power and thus can influence agents' reserve proposals or negotiation strategies; 2) agents can dynamically enter or leave the market (or negotiation); and 3) agents have incomplete information about the other entities.

Consider the following scenario in an above specified dynamic environment. An agent receives several proposals at round t and the agent has to make a decision on whether to accept the best proposal it has received as it's not sure that it can receive a better proposal in the future negotiation. This deliberation is indeed needed as agents may dynamically leave or enter the market. If the agent accepts the best proposal at round t, it may loose the opportunity to get a better negotiation result in the future. If it rejects the best proposal at round t, the agent has to face a risk of reaching a worse agreement (even no agreement) in the future negotiation.

An agent who tries to maximize profit by reaching an agreement at the most favorable prices possible (the highest price for the seller or the lowest price for the buyer) needs to reason about all its trading partners. One reasonable strategy for a negotiation agent would be to model what each individual agent thinks and will do, and use these models to figure out its best strategy. For example, in game theory, agents build a model of each other's possible moves and payoffs to find out their best moves (e.g., equilibrium strategies). However, modelling each of the other agents is often impossible or impractical in decision problems involving a large number of evolving participants as the result of agents' incomplete information and market uncertainty. Instead, we have developed a Markov chain stochastic modelling technique as part of an agent's negotiation strategy to make a decision on when to complete negotiation, which we call the MCDM (Markov chain based decision making) strategy. An agent with the MCDM strategy can take into account the dynamics and resulting uncertainties of the negotiation process using stochastic modelling of the negotiation process.

The remainder of this chapter is organized as follows. Section 2 presents the negotiation model and the MCDM strategy algorithm. In Section 3, we will explain the details of the Markov chain model. In section 4, we examine the performance of the MCDM strategy through experimentation. Section 5 summarizes related work. In the final section, some conclusions are presented and ideas for future work are outlined.

2 The Negotiation Model

For ease of analysis, this work focuses on single-issue (single attribute, e.g., price-only) negotiation rather than multiple-issue negotiation (we leave multiple-issue negotiation, which is more complex and challenging than a single-issue negotiation [6], for future research). This section presents 1) the negotiation protocol, 2) agents' concession strategy, 3) an agent's decision making problem in negotiation, and 4) the MCDM strategy algorithm.

2.1 Negotiation Protocol

To set the stage for specifying the negotiation model, some assumptions are given as follows:

1. Agents have incomplete information about other agents.
2. An agent concurrently negotiates with its trading partners.
3. Competition isn't considered.
4. Agents do not form coalitions.

Assumption 1 is intuitive because in practice, agents have private information, and for strategic reasons, they do not reveal their strategies, constraints (e.g., deadline, and reserve price), or preferences. In [14, p.54], it was noted that the strategy of a negotiation agent corresponds to its internal program, and extracting the true internal decision process would be difficult. Outside options can exist concurrently with a negotiation thread, or come sequentially in the future. A concurrently existing outside option is a negotiation thread that the negotiator is involved in simultaneously with another thread. Generally, a buyer gets more desirable negotiation outcomes when it negotiates concurrently with all the sellers in competitive situations in which there are information uncertainty and deadlines [11]. In this chapter, we assume that an agent negotiates concurrently with its trading partners (Assumption 2). As negotiation is a private behavior (only known to the two negotiators at most time), Assumption 3 is a plausible assumption. When buyers and sellers are allowed to group together to exploit the benefit of grouping, analysis of agents' strategies as well as market equilibrium become more complex [17]. At present stage, the design of negotiation agents does not consider coalition formation (Assumption 4).

Negotiation Set: As this work focuses on single-issue negotiation, the negotiation set or agreement set Φ, which represents the space of possible deals or proposals (see [14, p.34]) for an agent A is $[IP_A, RP_A]$ (IP_A and RP_A are, respectively, the initial and reserve prices of A). Let D be the event in which A fails to reach an agreement with its trading partner. Consequently, the utility function of A is defined as $U^A : \{\Phi\} \cup D \rightarrow [0,1]$ such that $U^A(D) = 0$ and for all $P_j^A \in [IP_A, RP_A]$, $U^A(P_j^A) > U^A(D)$. In the absence of an agreement, an agent receives a utility of zero, and it prefers all $U^A(P_j^A)$ to $U^A(D)$. If A is a buyer agent, then $U^A(P_1^A) > U^A(P_2^A)$ given that $P_1^A < P_2^A$, and $U^A(P_1^A) < U^A(P_2^A)$ for a seller agent, if $P_1^A > P_2^A$.

Negotiation Protocol: Negotiation proceeds in a series of rounds as follows. At round $t = 0$, the e-market opens. At any round, some new agents enter the market randomly. Negotiation begins when there are at least two agents of the opposite type (i.e., one buyer and one seller). In the first round of trading, an agent proposes a deal from its negotiation set. This work adopts the alternating offers protocol (see [15, p.100]) so that a pair of buyer and seller agents negotiates by making proposals in alternate rounds. Many buyer-seller pairs can negotiate deals simultaneously. If no agreement is reached, negotiation proceeds to another round. Negotiation between two agents terminates i) when an agreement is reached or ii) with a conflict when one of the two agents' deadline is reached.

2.2 Concession Strategy

The main goal of negotiation is to move towards and explore potential agreements within the common area of interest in order to find the most satisfactory agreement for the parties. In round $0 < t < min(T^B, T^{S_i})$ where T^B and T^{S_i} are, respectively, the deadlines of a buyer B and a seller S_i, if the proposal $P_{t-1}^{S_i \to B}$ of its trading partner at round $t-1$ isn't acceptable to the agent B, the agent B may make a concession to S_i at round t as reaching an acceptable agreement is always better than failing to reach an agreement.

Since a bargaining negotiation is fundamentally time-dependent [5], here we assume that both the two agents utilize a time dependent strategy while making a concession. During negotiation, penalty is incurred by one side or the other with the passage of time and each agent faces firm deadline.

Whereas deadline puts negotiators under pressure, they have different time preferences (e.g., negotiators with different time preferences may have different concession rates with respect to time). For instance, an agent may prefer to concede less rapidly in the early rounds of negotiation and more rapidly as its deadline approaches. The proposal of an agent A to its trading partner $\hat{A}$ at round t ($t < T^A$) is modelled as a time dependent function as follows:

$$P_t^{A \to \hat{A}} = IP^A - \phi^A(t) \times (IP^A - RP^A) \quad (1)$$

where t is current trading time.

The time dependent concession strategy in this work is used to decide the amount of concession in the price of a commodity. The time-dependent function $\phi^A(t)$ is determined with respect to time preference λ^A and deadline T^A (where $\lambda^A \geq 0$, and $T^A > 0$ is finite) and is given as

$$\phi^A(t) = (t/T^A)^{\lambda^A} \quad (2)$$

Although agents have infinitely many strategies with respect to remaining trading time (one for each value of λ^A), they can be classified as follows [17]:

1. *Linear (LN)*: $\lambda^A = 1$ and $\phi^A(t) = (t/T^A)$. At any round t, an agent makes a constant rate of concession $\phi^A(t-1) - \phi^A(t) = -1/T^A$.

2. *Conciliatory (CC)*: $\phi^A(t) = (t/T^A)^{\lambda^A}$, where $0 < \lambda^A < 1$. An agent makes larger concessions in the early trading rounds and smaller concessions at the later stage.

3. *Conservative (CS)*: $\phi^A(t) = (t/T^A)^{\lambda^A}$, where $1 < \lambda^A < \infty$. An agent makes smaller concessions in early rounds and larger concessions in later rounds.

4. "*Sit-and-wait (SW)*": In a bilateral negotiation, when both outside option and competition are not considered, if $\lambda = \infty$, then an agent adopts a "sit-and-wait" strategy. We can find that

$$\phi^A(t) = (t/T^A)^{\infty} = 0, \quad 0 \leq t < T^A$$

$$\phi^A(T^A) = (T^A/T^A)^{\infty} = 1, \qquad t = T^A$$

It is reminded that T^A is constant. For $t < T^A$, it follows that $(t/T^A)^\infty = 0$, and $P_t^{A\rightarrow\hat{A}} = IP^A$. When $t = T^A$, it follows that $(T^A/T^A)^\infty = 1$, and $P_{T^A}^{A\rightarrow\hat{A}} = RP^A$. Let $\triangle_t^A$ and $\triangle_{T^A}^A$ be the amounts of concession at $t < T^A$ and T^A, respectively. Before the deadline, an agent does not make any concession but "waits" for its trading partner to concede, since $\triangle_t^A = P_{t-1}^{A\rightarrow\hat{A}} - P_t^{A\rightarrow\hat{A}} = 0$ ($0 \leq t < T^A$). It only concedes at its deadline with $\triangle_{T^A}^A = P_{T^A-1}^{A\rightarrow\hat{A}} - P_{T^A}^{A\rightarrow\hat{A}} = IP^A - RP^A$.

The case when $\lambda = 0$ is not considered here because it represents a situation when no negotiation is needed.

Although there are four kinds of strategies for negotiation agents, in this chapter, we assume that all the negotiation agents take the "sit-and-wait" strategy since Sim [17, Proposition 3.1] has proved that the "sit-and-wait" strategy is the dominant strategy for an agent using time-dependent strategy, regardless of the strategy that its trading partner adopts. Therefore, a trading partner S_i of the buyer B will propose its optimal price IP^{S_i} at round $0 \leq j < T^{S_i} - 1$ and will propose RP^{S_i} at round $T^{S_i} - 1$. This assumption is consistent with an intuition noted by Raiffa (see [13, p.78]) that very often the strategic essence of a negotiation is merely a "waiting game", and a negotiator who is willing to wait, to probe more patiently, and to appear less eager for a settlement will be more successful. Although, in a multilateral negotiation, a negotiator that waits too long faces higher risk of losing a deal due to competition, it is shown in [16] that in a bilateral negotiation with deadlines, the only sequential equilibrium outcome is where each agent waits until the first deadline before accepting the proposal of its opponent.

Rational agents want to gain more through participating negotiation. Therefore, a buyer may let a very low price be its initial proposal and, in contrast, a seller may let a very high price be its initial proposal. For ease of analysis, this chapter assumes that an agent's optimal price isn't acceptable to its trading partners, i.e., for the negotiation pair B and S_i, $IP^{S_i} > RP^B$ and $IP^B < RP^{S_i}$. It can be found that the agreement can only be RP^B or RP^{S_i} (*Proposition 1*).

Proposition 1 *When a buyer B negotiates with a seller S_i and each agent's optimal price isn't acceptable to its trading partner, the agreement can only be RP^B or RP^{S_i} at round $min(T^B - 1, T^{S_i} - 1)$.*

Proof: Using the "sit-and-wait" strategy, each agent will propose its initial price before its deadline and will propose its reserve price at its deadline. Therefore, an agent only makes two different proposes: its initial price and its reserve price. Given the fact that each agent's initial price isn't acceptable to its trading partner, the possible agreement can only be one of their reserve prices, i.e., RP^B or RP^{S_i} for the negotiation between B and S_i at the shorter deadline $min(T^B - 1, T^{S_i} - 1)$ of the two agents. □

2.3 An Agent's Decision Problem

Usually a negotiator can face more than one candidate to reach an agreement, although only one agreement with a single candidate is allowed. These candidates become outside options with respect to each other for the negotiator. The outside options contribute to the environment of the negotiation with a candidate. Existence of outside options is typical in matching markets, and also common in commodity and service markets.

We call the negotiation between a negotiator and one of the trading partners a negotiation thread. Modelling the outside options and understanding the interaction between outside options and a negotiation thread is an essential aspect to design an effective negotiation strategy in the environment with outside options. The analysis of the decision model is presented from a buyer B's perspective while it negotiates with a set of sellers $S = \{S_1, S_2, \dots, S_n\}$ where n is the number of sellers. A similar model can be built from a seller's perspective. Note that n may dynamically change as we assume that agents dynamically enter or leave the market during negotiation. Let the number of trading partners of B at round t be n_t ($0 \leq t < T^B$).

The buyer B's utility function is given as:

$$U^B(P_j^B) = \frac{RP_B - P_j^B}{RP_B - IP_B} \tag{3}$$

where $P_j^B \in [IP_B, RP_B]$.

Now we give an analysis of the decision problem in negotiation. At round t ($0 \leq t < T^B$), the buyer B negotiates with n_t sellers and receives n_t proposals $P_t^{S \to B} = \{P_t^{S_i \to B} | i = 1, 2, \dots, n_t\}$ from its trading partners after it sent proposals to these trading partners at round t. Let the best proposal from B's perspective in $P_t^{S \to B}$ be $P_t^{S_j \to B}$, i.e., for any proposal $P_t^{S_i \to B} \in P_t^{S \to B}$, it follows that $U^B(P_t^{S_j \to B}) \geq U^B(P_t^{S_i \to B})$.[1] If $t = T^B - 1$, the buyer B has to accept the proposal $P_t^{S_j \to B}$ as it has no more time to bargain. If $t < T^B - 1$, the buyer has to make a decision on whether to accept the proposal $P_t^{S_j \to B}$. If it accepts the proposal $P_t^{S_j \to B}$, it will receive a payoff of $U^B(P_t^{S_j \to B})$ with certainty. If it rejects the proposal, the negotiation proceeds to the next round and the buyer B may 1) reach a better agreement than $P_t^{S_j \to B}$ if there is good chance in the future, 2) reach a worse agreement than $P_t^{S_j \to B}$ if there is no good chance in the future (e.g., the seller S_j may quit negotiation in the next round), or 3) be subjected to a conflict utility $U^B(D) = c^B = 0$ if there is no chance for the buyer to reach an agreement. c^B is the worst possible utility for B. Let $EU_{>t}^B$ be the buyer B's expected payoff it can receive in the future negotiation if it rejects the best proposal $P_t^{S_j \to B}$. B will find that it is advantageous to reject $P_t^{S_j \to B}$ at round t only if

$$U^B(P_t^{S_j \to B}) < EU_{>t}^B \tag{4}$$

Otherwise, it will accept $P_t^{S_j \to B}$ at round t.

Therefore, the decision problem for the buyer B at round t is to get the expected payoff $EU_{>t}^B$ in the future negotiation (after t). The agent B can use its beliefs about the market and its trading partners to compute the expected utility it can gain in the future negotiation. In this work, we compute the value of $EU_{>t}^B$ by use of a stochastic model in which the negotiation process is modelled by a Markov chain.

[1]There may be more than one proposal in $P_t^{S \to B}$ that generates the best utility for the buyer B. Without loss of generality, assume that each proposal generates different utility.

Function: The MCDM Strategy Algorithm.
Input: The agent B's beliefs about the market and its trading partners
Output: The value of $EU^B_{>t}$

Let $EU^B_{>t} = 0$.

Begin
For $t < t' < T^B$
Build a Markov chain from t to t' and compute the transition probabilities and rewards;
Compute the expected utility $EU^B_{t'}$ the agent B gains when it ends negotiation at round t';
If $EU^B_{t'} > EU^B_{>t}$ **Then**
$EU^B_{>t} = EU^B_{t'}$
End-if
End-For
Return $EU^B_{>t}$
End

Figure 1: The MCDM strategy algorithm.

2.4 The MCDM Strategy Algorithm

After rejecting the proposal $P_t^{S_j \to B}$ at round t, the buyer B can still reach an agreement with a seller at round t', $t < t' < T^B$. Let $EU^B_{t'}$ be the expected utility the agent B gets when it reaches an agreement at round t'. The value of $EU^B_{>t}$ is given by:

$$EU^B_{>t} = max_{t'} EU^B_{t'} \tag{5}$$

Fig.1 describes the MCDM strategy. As $EU^B_{>t} = max_{t'} EU^B_{t'}$, the buyer computes $EU^B_{t'}$ ($t < t' < T^B$) respectively and lets the highest value of $EU^B_{t'}$ be $EU^B_{>t}$. While computing the value of $EU^B_{t'}$, the buyer first builds a Markov chain from t to t' and computes the transition probabilities and rewards.

The MCDM strategy is a heuristic strategy. It models the dynamics of negotiation stochastically using a Markov chain (MC), assuming that the negotiation behaves as a random process. The MC model captures the variables that influence the agent's utility values and the uncertainties associated with them. The MC model takes those variables into account in the MC states and the transition probabilities. The details of the Markov chain model will be given in Section 3.

The MCDM strategy in Fig.1 is for the use of computing $EU^B_{>t}$ when the buyer makes a deliberation at round t ($0 \leq t < T^B$). If the buyer rejects $P_t^{S_j \to B}$ after the deliberation, the buyer still has to make the deliberation at round $t+1$ as it needs to compute $EU^B_{>t+1}$. Therefore, the buyer has to run the MCDM strategy algorithm in

Fig.1 for at most $T^B - 1$ times.

The complexity of the MCDM strategy algorithm is $O\big((T^B - t)\frac{|\Omega|^{(T^B-t)}-1}{|\Omega|-1}\big)$, where Ω is the set of different composition of sellers. Computing the expected utility $EU^B_{t'}$ takes $\frac{|\Omega|^{(t'-t)}-1}{|\Omega|-1}$ time (see Section 3). As the buyer needs to compute a set of values $\{EU^B_{t'}|t < t' < T^B\}$, the whole MCDM strategy algorithm will take $O\big((T^B - t)\frac{|\Omega|^{(T^B-t)}-1}{|\Omega|-1}\big)$ time. As the buyer has to run the MCDM strategy algorithm for at most $T^B - 1$ times, it will take $O\big((T^B - 1)T^B\frac{|\Omega|^{T^B}-1}{|\Omega|-1}\big)$ time to complete the negotiation.

3 The Markov Chain Model

This section presents the Markov chain model for computing the expected utility $EU^B_{t'}$ if the buyer B completes negotiation at round t'. This model is used for the buyer B's decision making at round t. In Section 3.1, we describe the variables captured in the MC model. In Section 3.2, how to model the negotiation process by the model is discussed. How to compute the transition probabilities between the MC states is described in Section 3.3. Given the MC model, the computation of the expected utility $EU^B_{t'}$ is discussed in Section 3.4.

3.1 Variables Captured in the MC Model

To get the expected payoff $EU^B_{>t}$ in the future negotiation, a MCDM strategy buyer needs to capture in its MC model the variables that influence the expected utility value. We divide those variables into three groups, as shown in Table 1. These variables define the buyer B's belief about the negotiation environments (including the market, its trading partners and itself). The variables in Table 1 are the information available in most negotiation scenarios, and we choose to capture all of them in the MC model.

The variable in the first group captures information about outside options. Following a usual way of modelling uncertain arrivals, we assume the arrival of outside options follows a probability distribution $P^O(n, t'')$ (in the experiments, we assume that the arrival of outside options follows a Poisson process), where $P^O(n, t'')$ denotes the probability that there will be n new sellers at round t''. This arrival probability $P^O(n, t'')$ together with its belief about the trading partners allows the buyer to forecast the number as well as the quality of the outside options arriving during the rest of the negotiation horizon.

The variables in the second group capture information about the trading partners of the buyer B. n_t is the number of trading partners at round t when it needs to make a decision on whether to complete negotiation. A negotiation has twosided incomplete information: both negotiation parties do not know the reservation price and the deadline of each other. Assume the buyer has an estimation of the reservation price of a seller, and the estimation is characterized by a probability distribution $F^R(x)$, where $F^R(x)$ denotes the probability that the reservation price of a seller is no higher than x. $F^R(x)$ is identical and independent across all sellers. This probability distribution is called the prior belief of the buyer. Similarly, a buyer also has an estimation of the deadline of a

Table 1: The buyer B's beliefs about the market and its trading partners

Information about the market	$P^O(n,t)$	Probability distribution of the arrival of outside options
Information about trading partners	n_t	Number of standing seller offers
	$F^R(x)$	Probability distribution of sellers' reserve prices
	$F^D(x)$	Probability distribution of sellers' deadlines
Information about self	RP^B	Its offer price
	U^B	Its utility function

seller, and the estimation is characterized by a probability distribution $F^D(x)$, where $F^D(x)$ denotes the probability that the deadline of a seller is no higher than x.

Finally, in addition to the information about the negotiation, the buyer needs information about itself – its initial price, its reserve price and its utility function.

3.2 Modelling the Negotiation Process

Using the information about the current status of the negotiation and beliefs about the negotiation environments, the buyer can determine the set of possible states, including the initial state, middle states and final states.

A negotiation state represents the status of negotiation. Therefore, a negotiation state should consist of two components: the set of trading partners and the corresponding negotiation time. Let Ω be the set of possible composition of trading partners. For example, $\Omega_i =< S_1 S_2 >$ represents that the buyer B is negotiating with two sellers S_1 and S_2. Each state can be represented as a pair (Ω_i, t'') which denotes that the buyer is negotiating with the set of agents in Ω_i at round t''. For convenience, in this work, a state (Ω_i, t'') is represented as $\Omega_i^{t''}$

Initial states: The initial state while building the Markov chain for getting the expected payoff $EU_{t'}^B$ is the set of sellers the buyer negotiates at round t. Let the initial state at time t be $\Omega_{Initial}^t$.

Contrast to the initial state, there are also middle states and final states. A state $\Omega_i^{t''}$ ($t < t'' < t'$) is called a middle state and a state $\Omega_i^{t'}$ is called a final state.

From the initial state, the MC models how the negotiation will proceed. As sellers may dynamically enter or leave the market, an initial state at round t'' may arrive at different states at round $t''+1$. Assuming that at most one trading partner leaves negotiation and at most one new trading partner enters negotiation, the negotiation can go to any of the following states from the $\Omega_i^{t''} =< S_1 S_2 >$ state:

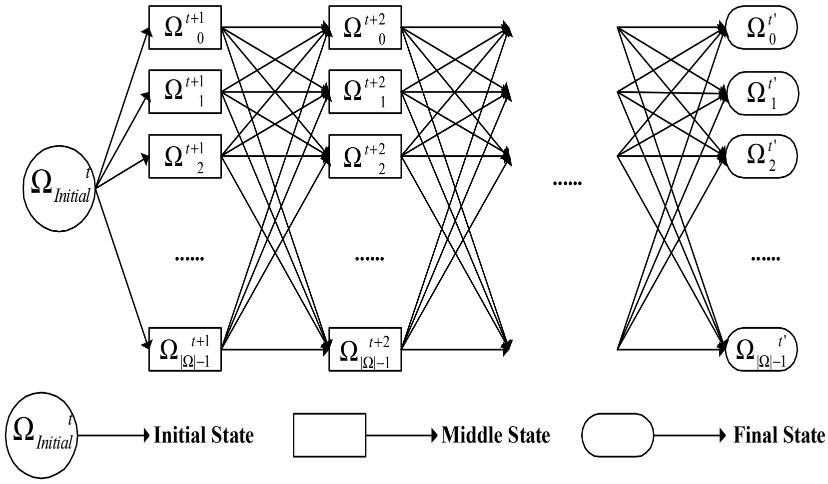

Figure 2: The MC model of the negotiation.

- $< S_1S_2 >$: No trading partner quits negotiation and no new trading partner enters negotiation.
- $< S_1 >$: S_2 quits negotiation and no new trading partner enters negotiation.
- $< S_2 >$: S_1 quits negotiation and no new trading partner enters negotiation.
- $< S_1S_3 >$: S_2 quits negotiation and a new trading partner S_3 enters negotiation.
- $< S_2S_3 >$: S_1 quits negotiation and a new trading partner S_3 enters negotiation.
- $< S_1S_2S_3 >$: No trading partner quits negotiation and a new trading partner S_3 enters negotiation.

All the remaining state transitions can be built in a similar way, and an example of a MC model is shown in Figure 2. From the initial state $\Omega^t_{Initial}$ of the negotiation, the process transitions to a set of possible middle states $\Omega^{t''}_i$, $t < t'' < t'$, and then to other middle states, and so on, until it goes to the final states $\Omega^{t'}_i$. There are $t' - t + 1$ layers of states in the MC model for computing the value of $EU^B_{t'}$, $t < t' < T^B$.

3.3 Computation of the Transition Probabilities

To complete the MC model, the buyer needs to compute the transition probabilities between the MC states. Note that although we describe building the MC model and computing its transition probabilities separately, they are in fact a one-step process: the buyer B computes the transition probabilities while building the MC model.

We use the transition from the state $\Omega^{t''}_i =< S_1S_2 >$ at time t'' ($t \leq t'' < t'$) to the $\Omega^{t''+1}_j =< S_1S_3 >$ state at time $t''+1$ as our example to illustrate the process. The following events result in the transition from the state $\Omega^{t''}_i$ to $\Omega^{t''+1}_j$:

I_1: The seller S_1 still doesn't quit negotiation.
I_2: The seller S_2 quits negotiation.
I_3: Only one new seller S_3 enters negotiation.

Let transition probability $P_{ij}^{t''}$ be the conditional probability that the process in state $\Omega_i^{t''}$ makes a transition into state $\Omega_j^{t''+1}$.

The events I_1, I_2 and I_3 result in the occurrence of the state transition. Therefore, the transition probability from the state $< S_1S_2 >$ at time t'' to the state $< S_1S_3 >$ at time $t'' + 1$, for example, is given as:

$$P_{ij}^{t''} = P\{\Omega_j^{t''+1} =< S_1S_3 > |\Omega_i^{t''} =< S_1S_2 >\} = P(I_1)P(I_2)P(I_3) \tag{3.1}$$

Where $P(I_1)$, $P(I_2)$ and $P(I_3)$ are, respectively, the probabilities of the occurrence of the events I_1, I_2 and I_3.

$P(I_1)$ is the probability that the seller S_1's deadline is longer than $t'' + 1 - T_{bg}(S_1)$ given that the seller S_1's deadline is longer than $t'' - T_{bg}(S_1)$, where $0 \leq T_{bg}(S_1) < T^B$ is the time when the seller S_1 begins to negotiate with the buyer B. Therefore,

$$P(I_1) = \frac{1 - F_{S_1}^D(t'' + 1 - T_{bg}(S_1))}{1 - F_{S_1}^D(t'' - T_{bg}(S_1))} \tag{3.2}$$

where $F_{S_1}^D(t'' + 1 - T_{bg}(S_1))$ is the probability that the seller S_1's deadline is no longer than $t'' + 1 - T_{bg}(S_1)$.

$P(I_2)$ is the probability that the seller S_2's deadline is $t'' + 1 - T_{bg}(S_2)$ given that the seller S_2's deadline is longer than $t'' - T_{bg}(S_2)$, where $0 \leq T_{bg}(S_2) < T^B$ is the time when the seller S_2 begins to negotiate with the buyer B. Therefore,

$$P(I_2) = \frac{F_{S_2}^D(t'' + 1 - T_{bg}(S_2)) - F_{S_2}^D(t'' - T_{bg}(S_2))}{1 - F_{S_2}^D(t'' - T_{bg}(S_2))} \tag{3.3}$$

$P(I_3)$ is the probability that only one new trading partner enters negotiation at round $t'' + 1$. Therefore,

$$P(I_3) = P^O(1, t'' + 1) \tag{3.4}$$

Similarly, the transition probability of any other transition process can be computed in the similar way.

3.4 Computing the Expected Utility Value

This section explains in detail how to compute the utility value $EU_{t'}^B$ for a given Markov chain Model.

The reward of a state represents the expected utility the buyer agent can gain from that state. Let the reward of a state $\Omega_i^{t''}$ be $R(\Omega_i^{t''})$. For any state $\Omega_i^{t''}$, $t < t'' < t'$, it follows that $R(\Omega_i^{t''}) = 0$ as the buyer will reject all the proposals before time t'. For a state $\Omega_i^{t'}$, $R(\Omega_i^{t'})$ depends on the quality of all the sellers' proposals. Take the state

$\Omega_i^{t'} =< S_1 S_2 S_3 >$ as an example. The reward of the state depends on the quality of the proposals of the sellers S_1 , S_2 and S_3. As 1) the agreement between two negotiation agents with the “sit-and-wait” strategy can only be one of their reserve prices and 2) each seller will propose its reserve price at its deadline, the expected utility the buyer can gain depends on 1) the probability that the seller S_1's deadline is $t' + 1 - T_{bg}(S_1)$ and 2) the expected reserve price of the seller S_1. Therefore, the expected utility the buyer can gain from the seller S_1 in $\Omega_i^{t'}$ is:

$$R(\Omega_i^{t'}, S_1) = P_{S_1}^D\Big(t' + 1 - T_{bg}(S_1)\Big)U^B\Big(\int_0^\infty f_{S_1}^R(x)dx\Big)$$

where $P_{S_1}^D(t' + 1 - T_{bg}) = \Big(F_{S_1}^D\big(t' + 1 - T_{bg}(S_1)\big) - F_{S_1}^D\big(t' - T_{bg}(S_1)\big)\Big)/\Big(1 - F_{S_1}^D\big(t' - T_{bg}(S_1)\big)\Big)$ is the probability that the seller S_1's deadline is $t' + 1 - T_{bg}(S_1)$ and $\int_0^\infty f_{S_1}^R(x)dx$ is the expected reserve price of the seller S_1 in which $f_{S_1}^R(x)$ is the probability density function of the expected reserve price of the seller S_1.

Similarly, let the expected utilities the buyer can gain from the sellers S_2 and S_3 be $R(\Omega_i^{t'}, S_2)$ and $R(\Omega_i^{t'}, S_3)$ respectively. The reward of a state $\Omega_i^{t'}$ will be the maximum of the three expected utilities, i.e., $R(\Omega_i^{t'}) = max\big(R(\Omega_i^{t'}, S_1), R(\Omega_i^{t'}, S_2), R(\Omega_i^{t'}, S_3)\big)$.

Formally, the reward $R(\Omega_i^{t'})$ of the final state $\Omega_i^{t'}$ can be defined as:

$$R(\Omega_i^{t'}) = max_{S_j \in \Omega_i^{t'}} P_{S_j}^D\Big(t' + 1 - T_{bg}(S_j)\Big)U^B\Big(\int_0^\infty f_{S_j}^R(x)dx\Big) \tag{3.5}$$

where $P_{S_j}^D\big(t' + 1 - T_{bg}(S_j)\big) = \Big(F_{S_j}^D\big(t' + 1 - T_{bg}(S_j)\big) - F_{S_j}^D\big(t' - T_{bg}(S_j)\big)\Big)/\Big(1 - F_{S_j}^D\big(t' - T_{bg}(S_j)\big)\Big)$.

Given the reward of each state, the transition probability from one state to another state and the initial state $\Omega_{Initial}^t$, the expected utility $EU_{t'}^B$ if the buyer completes negotiation at round t' can be given as (assume that $\Omega_{Initial}^t = \Omega_j^t$):

$$EU_{t'}^B = \sum_{\Omega_i \in \Omega} P_{ji}^t V(\Omega_i^{t+1}) \tag{3.6}$$

where $V(\Omega_i^{t''})$, $t < t'' \leq t'$ is given as:

$$V(\Omega_i^{t''}) = \begin{cases} R(\Omega_i^{t''}) & \text{if } t'' = t' \\ R(\Omega_i^{t''}) + \sum_{\Omega_j \in \Omega} P_{ij}^{t''} V(\Omega_j^{t''+1}) & \text{if } t'' < t' \end{cases} \tag{3.7}$$

As $R(\Omega_i^{t''}) = 0$ when $t < t'' < t'$, (3.7) can be rewritten as

$$V(\Omega_i^{t''}) = \begin{cases} R(\Omega_i^{t''}) & \text{if } t'' = t' \\ \sum_{\Omega_j \in \Omega} P_{ij}^{t''} V(\Omega_j^{t''+1}) & \text{if } t'' < t' \end{cases} \tag{3.8}$$

It can be found from (3.6) and (3.8) that the time complexity of computing $EU_{t'}^B$ is $\frac{|\Omega|^{(t'-t)}-1}{|\Omega|-1}$.

Table 2: Input data sources

Input Data	Possible Values		
Market Density	*Sparse*	*Moderate*	*Dense*
N_{gen}	0.5	2	4
N_{gen}: The *average number* of *agents generated* per round			
Deadline	*Short*	*Moderate*	*Long*
T_{max}	$18-25$	$35-45$	$60-70$

4 Experimentation

4.1 Testbed

To realize the idea of the Markov chain based decision making model, a simulation testbed consists of a virtual e-Marketplace, a society of trading agents and a controller (manager) was implemented. The controller generates agents, randomly determines their parameters (e.g., their roles as buyers or sellers, initial proposals, reserve price, negotiation mechanisms, deadlines), and simulates the entrance of agents to the virtual e-Marketplace. Using the testbed a series of experiments were carried out in order to evaluate the effectiveness of the Markov chain based decision making model. In order to demonstrate the performance of the Markov chain based decision making model by comparison, the general decision strategy was also evaluated.

General strategy: An agent with the general decision strategy will reach an agreement while one of its trading partners's counter-proposals is better than its proposal, and accepts the best proposal of the proposals from its trading partners. Therefore, an agent with the general strategy will not make any deliberation during negotiation but accepts the best proposal at the last round. In contrast, an agent with the MCDM strategy will make a deliberation during each round and may finish negotiation before its deadline.

To evaluate the performance of the two decision strategies in a wide variety of test environments, agents are subject to different market densities and different deadlines (Table 2).

4.2 Experimental Settings

Both the two input parameters in Table 2 are generated randomly following a uniform distribution. Market density depends on average number of agents generated per round. The lifespan of an agent, i.e., its deadline, is randomly selected from $[15, 70]$. The range of $[15, 70]$ for deadline is adopted based on experimental tuning and agents' behaviors. In current experimental setting, it was found that: 1) for very short deadline (< 15), very few agents could complete deals, and 2) for deadlines > 70, there was little or no difference in performance of agents. Hence, for the purpose of experimentation, a deadline between the range of $18-25$ (respectively, $35-45$ and $60-70$) is considered as short (respectively, moderate and long).

This work assumes that the arrival of outside options and agents' deadlines follows a Poisson process. An agent (e.g., a buyer)'s beliefs about the market and its trading

Table 3: Performance Measure

Success Rate	$R_{success} = N_{success}/N_{total}$
Expected Utility	$U_{expected} = U_{success} \times R_{success} + U_{fail} \times (1 - R_{success}) = U_{success} \times R_{success}$
Average Negotiation Time	$R_{time} = \sum_{i=1}^{N_{total}} T_{end}^{i}/N_{total}$
N_{total}	*Total number of agents*
$N_{success}$	*No. of agents that reached consensus*
$U_{success}$	*Average utility of agents that reached consensus*
$U_{fail} = 0$	*Average utility of agents that didn't reach consensus*
T_{end}^{i}	*The time spent in negotiation by the agent i*

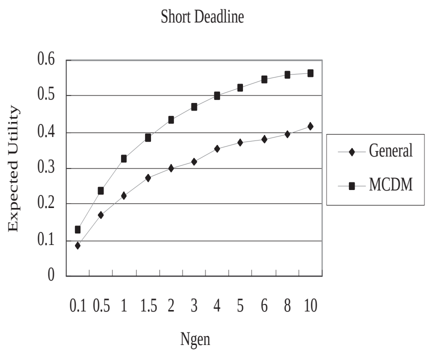

Figure 3: Expected utility and market dynamics.

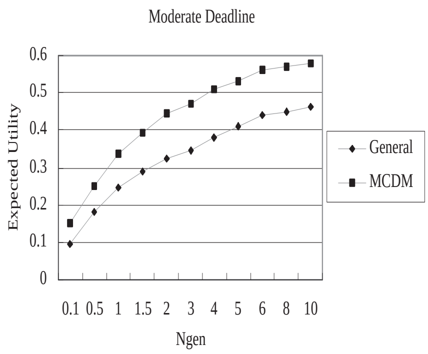

Figure 4: Expected utility and market dynamics.

partners are affected by the above two input data. For example, in dense market where there are more new negotiators, an agent will believe that the probability of arriving more new trading partners is higher than that in moderate market and sparse market.

4.3 Performance Measure

The performance measures include expected utility, success rate and average negotiation time (Table 3). Other than optimizing agents' utility, enhancing the success rate is also an important evaluation criteria for designing negotiation agents as pointed out in [4, 19]. Since negotiation outcomes are uncertain (i.e., there are two possibilities: eventually reaching a consensus or not reaching a consensus), it seems more prudent to use expected utility [2] (rather than average utility) as a performance measure since it takes into consideration the probability distribution over the two different outcomes [19]. Average negotiation time examines the average amount of time spent in negotiation.

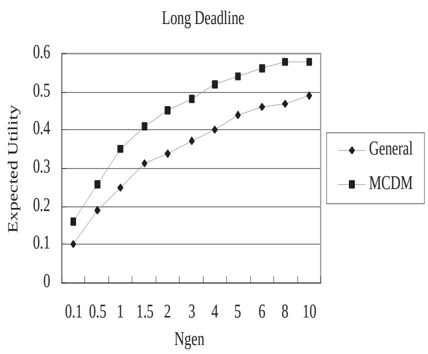

Figure 5: Expected utility and market dynamics.

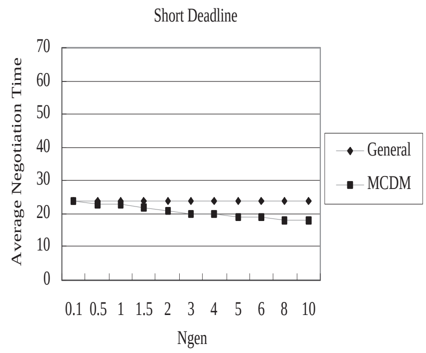

Figure 6: Average negotiation time and market dynamics.

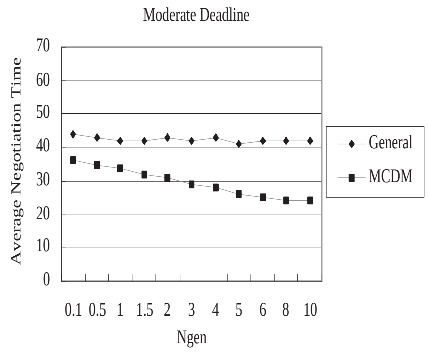

Figure 7: Average negotiation time and market dynamics.

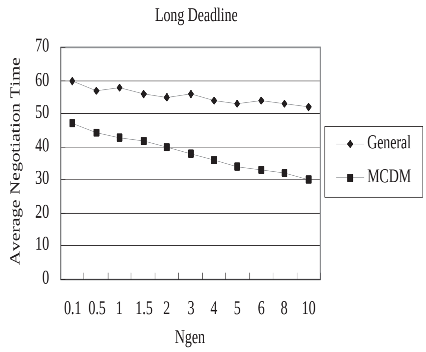

Figure 8: Average negotiation time and market dynamics.

4.4 Results

An extensive amount of stochastic simulations were carried out for all the combinations of market density (dense, moderately dense, sparse)and deadline (short, moderate, long). An extensive amount of experiments were carried out for all 9 (3×3) combinations of the input data and some representative results are presented in Figs. 3–11.

4.5 Observation 1

When both agents with the general strategy and agents with the MCDM strategy are subject to different market densities, agents with the MCDM strategy always achieved higher $U_{expected}$ than agents with the general strategy. For example, in Fig. 4, when the average number of new trading partners is 2 in each round, the expected utilities are 0.44 for agents with the MCDM strategy, and 0.33 for agents with the general strategy respectively.

When the market becomes more dynamic, the advantage of utilities of agents with the MCDM strategy over agents with the general strategy increases. This corresponds to the intuition that the potential of increasing agents' utilities increases in dynamic environments where agents can dynamically enter or leave negotiation. For example, in Fig. 4 where agents have moderate deadlines, the expected utilities are 0.25 for agents with the MCDM strategy, 0.18 for agents with the general strategy when $N_{gen} = 0.5$ respectively; 0.44 for agents with the MCDM strategy, 0.33 for agents with the general strategy when $N_{gen} = 2$ respectively; and 0.53 for agents with the MCDM strategy, 0.41 for agents with the general strategy when $N_{gen} = 5$ respectively.

With the increase of deadline, the advantage of utilities of agents with the MCDM strategy over agents with the general strategy decreases. This also corresponds to the intuition that the potential of increasing agents' utilities increases when deadlines become shorter, which will make the market more dynamic. In Figs. 3, 4 and 5, when the average number of new trading partners in each round is 5, the expected utilities are 0.52 for agents with the MCDM strategy, 0.37 for agents with the general strategy when agents have short deadlines respectively; 0.53 for agents with the MCDM strategy, 0.41 for agents with the general strategy when agents have moderate deadlines respectively; and 0.54 for agents with the MCDM strategy, 0.44 for agents with the general strategy when agents have long deadlines respectively.

4.6 Observation 2

Through the experimental results in Figs. 6-8, we can find that: with different levels of market dynamics, agents with the MCDM strategy always achieved lower R_{time} than agents with the general strategy. For example, in Fig. 7, when the average number of new trading partners in each round is 2, the average negotiation times are 31 for agents with the MCDM strategy, and 43 for agents with the general strategy respectively.

When the market becomes more dynamic, the advantage of negotiation time of agents with the MCDM strategy over agents with the general strategy increases. For example, in Fig. 7 where agents have moderate deadlines, the average negotiation times are 35 for agents with the MCDM strategy, 43 for agents with the general strategy when $N_{gen} = 0.5$ respectively; 31 for agents with the MCDM strategy, 43 for agents with the general strategy when $N_{gen} = 2$ respectively; and 25 for agents with the MCDM strategy, 42 for agents with the general strategy when $N_{gen} = 5$ respectively.

Through comparison of the negotiation results in Figs. 6, 7 and 8, we can also find that with the increase of deadline, the advantage of negotiation time of agents with the MCDM strategy over agents with the general strategy increases.

4.7 Observation 3

When both agents with the general strategy and agents with the MCDM strategy are subject to different market densities, agents with the MCDM strategy always achieved higher $R_{success}$ than agents with the general strategy. For example, in Fig. 10, when the average number of new trading partners is 2 in each round, the success rates are 0.38 for agents with the MCDM strategy, and 0.22 for agents with the general strategy respectively.

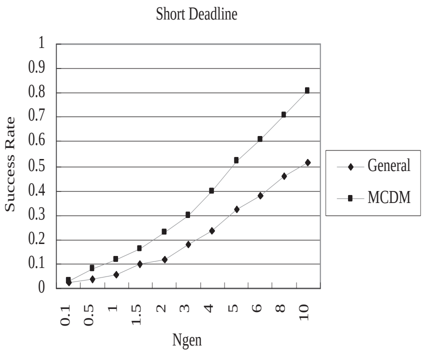

Figure 9: Success rate and market dynamics.

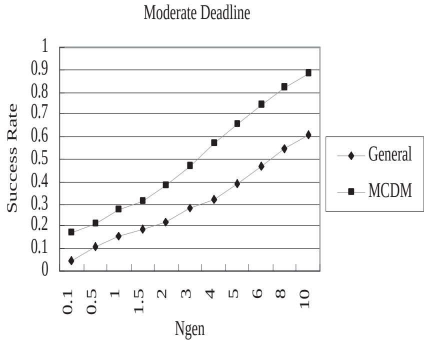

Figure 10: Success rate and market dynamics.

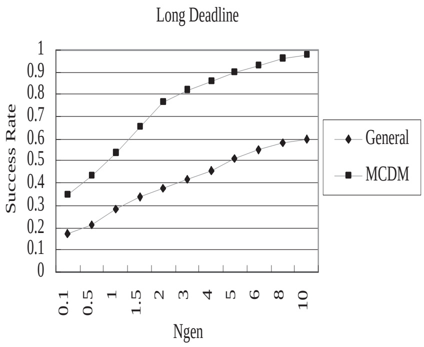

Figure 11: Success rate and market dynamics.

We can also find that, when the market becomes more dynamic, the advantage of success rates of agents with the MCDM strategy over agents with the general strategy increases. This corresponds to the intuition that the potential of increasing agents' chance of reaching agreements increases in dynamic environments where more agents can dynamically enter or leave negotiation. Moreover, with the increase of deadlines, the advantage of success rates of agents with the MCDM strategy over agents with the general strategy increases. This phenomenon corresponds to the intuition that the potential of increasing agents' success rate decreases under very extreme (or very adverse) trading conditions like short deadlines.

5 Related Work

The literature of automated negotiation and negotiation agents forms a very huge collection, and space limitations preclude introducing all of them here. For a survey on automated negotiation, see [4, 9]. The rest part of this section only introduces and discusses some important related work with respect to negotiation agents in complex and dynamic environment.

The environment that a negotiator is situated in greatly impacts the course of negotiation actions. Instead of focusing on analyzing the strategy equilibrium as a function of (the distribution of) valuations and historical information as in game theory, researchers in AI are interested in designing flexible and sophisticated negotiation agents in complex environments. Faratin et al. [3] devised a negotiation model that defines a range of strategies and tactics for generating proposals based on time, resource, and behaviors of negotiators, which are widely used in automated negotiation. However, they didn't consider the influence of outside options. Moreover, market dynamics were ignored in their work.

Nguyen and Jennings [10, 11] have developed and evaluated a heuristic model that enables an agent to participate in multiple, concurrent bilateral encounters in competitive situations in which there is information uncertainty and deadlines. The main findings through empirical evaluation include: 1) The time to complete the negotiation is less for the concurrent model than for the sequential one; 2) To realize the benefits of concurrent negotiation, the buyer agent's deadline must not be too short; 3) The final agreements reached by the concurrent model have, on average, higher or equal utility for the buyer than those of the sequential model; 4) Changing the strategy in response to the agent's assessment of the ongoing negotiation is equal or better than not doing so; 5) To improve the performance of the concurrent model, the analysis time should be moderately early but not too early; 6) The tougher the buyer negotiates the better the overall outcome it obtains. In [10, 11], multiple negotiation threads are assumed to be independent on one another. In contrast, in this work, outside options' influence is considered in an agent's decision making for a single negotiation thread.

Although strategies in [3] are based on time, resource, and behaviors of negotiators, other essential factors, such as competition, trading alternatives, and differences of negotiators are not considered. To take the important factors ignored in [3] into account while designing e-negotiation agents, Sim et al. [17–20] have designed and implemented a society of market-driven agents that make adjustable amounts of concession by reacting to market dynamics. Previous empirical results in [18] show that in general, market-driven agents achieve trading outcomes that correspond to intuitions in real-life trading. The market situations include trading opportunities, competition, remaining trading time, and eagerness. Multiple trading opportunities in the market can be regarded as outside options against each other for a negotiator. In their model the number of trading opportunities influences the aggregated probability of conflict, which determines the probability of completing a deal in the current negotiation cycle. With more trading opportunities, the probability of completing a deal is higher, and it follows that the negotiator's concession is smaller in the next cycle based on the spread decision function. This work complements their work by considering outside options' effect on negotiation results and agents can make decisions accordingly.

Li et al. [7, 8] present a model for bilateral contract negotiation that considers the uncertain and dynamic outside options. Outside options affect the negotiation strategies via their impact on the reserve price. The model is composed of three modules: single-threaded negotiation, synchronized multi-threaded negotiation, and dynamic multi-threaded negotiation. These three models embody increased sophistication and complexity. The single-threaded negotiation model provides negotiation strategies without specifically considering outside options. The model of synchronized multi-threaded negotiation builds on the single-threaded negotiation model and considers the presence of concurrently existing outside options. The model of dynamic multi-threaded negotiation expands the synchronized multi-threaded model by considering the uncertain outside options that may come dynamically in the future. The discrete time concurrent one-to-many negotiation model by Li et al. [7,8] assumes that an agent has information about the expected utility of an outside option but such information is difficult to get. In contrast, this work assumes that an agent only has information about the probability distribution of its trading partners' reserve price.

6 Conclusions

This research investigates an agent's decision making on when to complete negotiation in dynamic and complex negotiation environments. Unlike the existing general decision strategy in which an agent accepts the best proposal by its deadline, an agent using the MCDM strategy can make a decision on when to complete negotiation according to its beliefs about market dynamics and its trading partners, which will help with agents' utility optimization as validated in our experiments.

We have empirically evaluated the performance of the MCDM strategy. First, we have compared the MCDM strategy agents to agents using the general strategy under various environments. The results indicate that the MCDM strategy outperforms the general strategy. The simulation outcomes in Section3 4.5–4.6 suggest that MCDM strategy always gets results of higher utility (respectively, shorter average negotiation time and higher success rate) with different market densities and deadlines. In particular, the MCDM strategy agents perform well in dynamic environments. For example, when there are many new negotiators appear in each round, the advantages of the MCDM strategy over the general strategy is high.

In summary, the proposed MCDM strategy introduced in this research helps to agents' utility optimization and adaptation to market dynamics. The proposed strategy can be applied in open, dynamic, and complicated negotiation environments (such as, service oriented Grid, supply chain and workflow). However, this research makes no claim that our strategies are sufficiently accurate or powerful to solve all or most of the problems in automated negotiation. When agents with our negotiation strategies are engineered with some elements that model how negotiators in reality may behave, it is not the intention of this research to create an exact replica of negotiators in realistic markets.

Finally, a future agenda of this work is engineering behavior-based tactics [3] and learning techniques [21] into agents' decision making in automated negotiation. Another interesting extension is analyzing agents' decision making within the continuous time negotiation mechanism [1].

7 Acknowledgement

Early version of this paper was appeared in the Proceedings of the AAMAS–06 Workshop on Rational, Robust, and Secure Negotiations, pp.1–19. The authors would like to thank the anonymous referees and the guest editors of this special issue for their valuable comments and suggestions for improving this paper.

References

[1] B. An, K. M. Sim, L. G. Tang, S. Q. Li, and D. J. Cheng, “A continuous time negotiation mechanism for software agents,” *IEEE Trans. on Systems, Man and Cybernetics, Part B*, vol. 36, no. 6, Dec. 2006.

[2] P. K. Dutta, *Strategies and games: theory and practice*. Cambridge, MA: MIT Press, 1999.

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

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

[5] S. Kraus, “Automated negotiation and decision making in multi-agent environments,” *Multi-Agents systems and applications*, no. 2, 2001.

[6] G. Lai, C. Li, K. Sycara, and J. Giampapa, “Literature review on multi-attribute negotiations,” Robotics Institute, Carnegie Mellon University, Tech. Rep., Dec. 2004.

[7] C. Li, J. Giampapa, and K. Sycara, “Bilateral negotiation decisions with uncertain dynamic outside options,” in *Proc. of the First IEEE Int. Workshop on Electronic Contracting*, 2004, pp. 54–61.

[8] C. Li, J. A. Giampapa, and K. Sycara, “Bilateral negotiation decisions with uncertain dynamic outside options,” *IEEE Trans. on Systems, Man, and Cybernetics, Part C: Special issue on Game-theoretic Analysis and Stochastic Simulation of Negotiation Agents*, vol. 36, no. 1, pp. 31–44, Jan. 2006.

[9] A. R. Lomuscio, M. Wooldridge, and N. R. Jennings, “A classification scheme for negotiation in electronic commerce,” *Int. Jounal Group Decision Negotiation*, vol. 12, no. 1, pp. 31–56, 2003.

[10] T. D. Nguyen and N. R. Jennings, “Concurrent bi-lateral negotiation in agent systems,” in *Proc. of the fourth DEXA Workshop on e-Negotiations*, 2003, pp. 839–844.

[11] ——, “A heuristic model of concurrent bi-lateral negotiations in incomplete information settings,” in *Proc. of the eighteenth Int. Joint Conference on Artificial Intelligence*, 2003, pp. 1467–1469.

[12] D. Pruitt, *Negotiation behavior*. New York: Academic Press, 1981.

[13] H. Raiffa, *The art and science of negotiation*. Cambridge, MA: Harvard Univ. Press, 1982.

[14] J. Rosenschein and G. Zlotkin, *Rules of encounter: designing conventions for automated negotiation among computers*. Cambridge, MA: MIT Press, 1994.

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

[16] T. Sandholm and N. Vulkan, "Bargianing with deadlines," in *Proc. of AAAI99*, 1999, pp. 44–51.

[17] K. M. Sim, "Equilibria, prudent compromises, and the "waiting" game," *IEEE Trans. on Systems, Man and Cybernetics, Part B*, vol. 33, no. 3, pp. 712– 724, Jun. 2005.

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

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

[20] K. M. Sim and E. Wong, "Towards market-driven agents for electronic auction," *IEEE Trans. on Systems, Man and Cybernetics, Part A*, vol. 31, no. 6, pp. 474–484, Nov. 2001.

[21] D. Zeng and K. Sycara, "Bayesian learning in negotiation," *Int. Journal of Human Computer Systems*, vol. 48, no. 3-4, pp. 125–141, 1998.