Integration of Learning, Situational Power and Goal Constraints Into Time-Dependent Electronic Negotiation Agents

by

Wilson Wai Ho Mok

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Applied Science
in
Management Sciences

Waterloo, Ontario, Canada, 2002

© Wilson Wai Ho Mok 2002

Author's Declaration For Electronic Submission of a Thesis

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

In the past decade, electronic negotiation has become an important research topic in the field of information systems. A desirable goal of negotiation agents is to understand their owners' requirements, and to learn their opponents' behavior, thereby lessening the involvement of human beings. Studies on human negotiation bring out that several issues can affect a human's negotiation behavior, including learning an opponent's behavior, exerting power on an opponent, and setting an individual goal to improve the level of accomplishment. Research on incorporating these issues into negotiation agents is, however, still at an infancy state. We therefore take up this topic in this thesis.

Researchers have proposed many different negotiation agents that follow a preset behavior based on human models of negotiation. In this thesis, we consider one such model, known as the time-dependent-tactical model, which is used by human negotiators and in which the values of the negotiating issues are determined based on the time elapsed in the negotiation. A learning mechanism for this model might be beneficial, because this model is frequently used in electronic negotiation. Thus, we propose heuristic algorithms that estimate the parameters of an agent's time-dependent-tactical model, and that then react to the estimated parameters for achieving higher negotiation performance. Besides learning, we incorporate two other factors that have been found to affect a human negotiation outcome. These are situational power, which represents differences in negotiators' status based on market conditions, and goal constraints, which stand for the levels of accomplishment negotiators try to strive for. To validate the impacts of learning, situational power and goal constraints in electronic negotiation, we first present how to integrate these features into negotiation agents, and then conduct simulations. With 187,500 simulation runs, we observe that our learning algorithms are effective in improving both individual and dyadic negotiation performances. For the effects of situational power and goal constraints, we obtain congruent results between human and electronic negotiations. By incorporating learning into situational power and goal constraints, we achieve significant joint effect between learning and situational power as well as that between learning and goal constraints.

In summary, this thesis provides three primary contributions to the fields of information systems and electronic-commerce research. First, we have designed algorithms for learning an opponent's negotiation behavior. Second, our learning algorithms are found to be effective in improving negotiation performance. Third, we have shown how learning can be integrated with situational power and goal constraints, although this is not a major focus in this study. Finally, the agreement on the joint effects of learning, situational power and goal constraints between human and electronic negotiations suggests that our integrated design of the agent appears to be effective.

Acknowledgements

I would like to thank Dr. R. Sundarraj, who was my supervisor for this thesis. Without his support and encouragement, my contribution to this thesis would not have been possible. I really appreciate the effort and the time he put in for helping me to prepare a well-organized and well-written paper.

I would also like to say thanks to my two readers, Dr. E. M. Jewkes and Dr. Ji-Ye Mao, who would spend time on reading my thesis and providing feedback.

Table of Contents

List of Tables

List of Illustrations

1 Introduction

1.1 Motivation

Electronic agents are becoming an increasingly important research topic in the field of information systems. An electronic agent is a software program that represents its owner to perform specific tasks. To do this, it needs to be autonomous, reactive and proactive (Jennings, 2001). Agents come in different types. Information-filtering agents, motivated by the information explosion on the Internet, help users to gather and filter information about a chosen topic. Shopping agents assist users to shop online, by providing information, explanation, and recommendation to users who lack knowledge on their interested products (Murugesan, 1998). In addition, with the rapid development of electronic commerce, there are agents that help organizations and people in conducting negotiations (Ye *et al.*, 2001). The focus of this thesis is on negotiation agents.

Negotiation agents act on behalf of customers or sellers and negotiate with other agents on several issues such as the price and the delivery time of a product. The motivation for negotiation agents stems from the fact that human negotiation is slow and complex, and makes negotiations inefficient (Murugesan, 2000; Choi *et al.*, 2001). As a typical negotiation involves multiple issues, the process of attaining an optimal outcome for all parties is comparatively complex. Thus, achieving the optimal outcome becomes difficult for human negotiators (Raiffa, 1982, Pruitt, 1981), and human negotiators often end up with a sub-optimal outcome. This induces an increasing need for an electronic negotiation agent that can provide a higher level of information-processing capabilities (Lim & Benbasat, 1993).

Since negotiation agents are ultimately meant to complement and aid human negotiation, a desirable element of such an agent is to understand its owner's preferences, and to learn its opponents' negotiation behavior (Choi *et al.*, 2001). Motivated by this goal, researchers have proposed many different behavioral negotiation agents that attempt to follow preset models drawn from human negotiation. One of these models is time-dependent-tactical model, by which agents determine the values of negotiating issues based on the time elapsed in the negotiation. Previous research in human negotiation point out that several issues such as learning and psychological factors can affect a human's negotiation behavior. Psychologists have determined that a negotiator's demands can depend on the concession he/she expects from other parties (Pruitt, 1981; Raiffa, 1982). Thus, if agents can learn and understand the time-dependent negotiation behavior of their opponents, they might improve their negotiation performance by adjusting their own negotiation models. It is, however, the case that research on incorporating learning mechanisms into electronic negotiation is still at a preliminary stage.

In addition, psychologists have determined that situational power and goal constraints can affect a human's negotiation behavior. Situational power represents the differences in negotiators' status according to the market conditions. Negotiators who see themselves as powerful usually extract a large concession from their opponents, and make small concession of their own (Pruitt, 1981; Rubin and Brown, 1975). Goal constraints stand for the levels of accomplishment negotiators try to strive for. With more difficult goals, negotiators have higher motivation, and thus work harder to satisfy their goals (Locke *et al.*, 1981). Integration of these two elements has already been put forth in Sundarraj (2002), but the effects of these two factors on the level of agreement between human and electronic negotiations have not been validated yet. This validation is important because it evaluates the extent to which the integration mechanisms are useful in electronic negotiation tasks.

To close the gaps outlined above, we do the following.

1) We design algorithms for learning an opponent's negotiation behavior, with the time-dependent tactic as the basis.
2) We integrate situational power and goal constraints into negotiation agents. Since the fundamental idea of this integration stems from Sundarraj (2002), we treat this objective as a minor one. We adapt from his design and further modify the integration methodologies.
3) We conduct experiments to test the effectiveness of the learning algorithms, and the effects of situational power and goal constraints on electronic negotiation.

1.2 Overview

To fulfill the first objective given above, we develop learning algorithms to estimate the parameters of an agent's time-dependent-tactical model. Our algorithms consist of three phases: 1) Preprocessing Phase, 2) Construction Phase, and 3) Improvement Phase. In these phases, we progressively reduce the enumeration range for estimating the parameters, and fine-tune the parameters to fit the received offers. In the preprocessing phase, we use an agent's initial demand level to construct a third-degree Taylor's series approximation of the negotiation function. Based on the Taylor series, we estimate the agent's concession rate, and use it to determine the type of concession behavior. In the construction phase, we use the received offers, the estimated concession rates from the preprocessing phase, the time-dependent tactic equation, and its derivatives' expressions to find a reference set of values for the parameters of the time-dependent tactic. In the improvement phase, we carry out enumeration over a discretized range of parameter values based on the reference set of values from the construction phase. The estimated parameters are then used to adjust the agent's own tactic, so that higher negotiation performance can be achieved.

For our second objective, the situational power of an agent is defined as its perception about the preferences of potential buyers in the market for the product in question. Situational power is manipulated using the analytical hierarchy process (AHP) method. This use of AHP is similar to the application of AHP in benchmarking the performance of an organization (Forman & Gass, 2001). Specifically, by eliciting an agent's perception of the market preferences using pairwise comparisons of the product's attributes and then comparing these preferences with a neutral set of preferences, we can obtain a measure of an agent's perceived power. We then use this measure of power to adjust the agent's negotiation model. To manipulate goal, we allow users to specify threshold utility values that must be achieved before concluding the negotiations.

Finally, to test the effect of our manipulation on electronic negotiation, we conduct two sets of experiments. The first set consists of 33,750 negotiation simulations, with three levels of perceived power and the type of the agent (buyer or seller) as two factors. For the effect of goal constraints, we carry out 60,000 simulations, with four levels of goal and the agent type as the two factors. Both these experiments are run with and without learning, and statistically analyzed using a range of individual and dyadic performance measures from the literature.

The thesis is organized as follows. In Chapter 2, we review the models that have been used in electronic negotiations, and present how our work is different from these studies. In Chapter 3, we identify the concepts and the significance of time-dependent tactic, and present how agents can employ this tactic to negotiate. Then we discuss the concept of our heuristic algorithms for learning in Chapter 4, and in Chapter 5, we describe the experiment for assessing the extent to which the algorithms can estimate the parameters of time-dependent tactic. In Chapter 6, we introduce how to integrate situational power, goal constraints as well as learning into negotiation agents. Subsequently, in Chapter 7, we propose the hypotheses of the effects of situational power, goal constraints and learning on electronic negotiation, and address the experimental design. The results and discussion for the power-setting and goal-setting experiments are presented in Chapters 8 and 9, respectively. Chapter 10 concludes with a discussion of the major results and future works.

2 Literature Review

The rapid expansion of electronic commerce in the past decade has been the motivation for many studies focusing on automated negotiation agents. These agents can ensure speed and consistency, mitigate the effects of human error, and offer all-time availability for global e-business market (Murugesan, 2000). In this chapter, our purpose is to review previous studies on electronic negotiations, and present our contributions in the context of these studies.

First, we discuss the definition of negotiation in detail in §2.1. In §2.2, we review current research on electronic negotiation. Finally, we present our contributions in §2.3.

2.1 Definition of Negotiation

Negotiation is a complex iterative process, and is common in many different fields such as business and politics. It can simply represent the purchase of an item between a buyer and a seller in a business transaction, or a peace negotiation agreement between two opposing parties. According to Pruitt (1981), "*Negotiation is a process by which a joint decision is made by two or more parties. The parties first verbalize contradictory demands and then move toward agreement by a process of concession making or search for new alternatives*". Negotiation can also be defined as a process of searching a feasible alternative to resolve conflict between parties (Choi *et al.*, 2001). Basically, negotiation is a process to reach an agreement between two or more parties on certain terms such as price and quantity. One party proposes an offer, while another party evaluates and decides whether to accept or reject the offer. If the party rejects, it may propose a counter offer. This process is iterated until both parties compromise for an agreement, or either party terminates the negotiation.

Negotiation can be classified as distributive or integrative bargaining (Goh *et al.*, 2000; Teich *et al.*, 1999). In distributive bargaining, two parties bargain over a fixed pie. Both parties have high interest in the same issue, and thus create a high conflict. On the other hand, in integrative bargaining, two parties expand the pie when they have interests in different issues. Under this situation, conflict is reduced, and both individual and dyadic performances increase. One example to present the difference between distributive and integrative bargaining is the division of a glass of apple juice and a glass of orange juice between two persons. If person A likes to drink apple juice while person B likes to drink orange juice, each person receives a half glass of each type of juice in distributive bargaining. On the other hand, in integrative bargaining, person A chooses to have all apple juice while person B chooses to have all orange juice.

2.2 Research into Electronic Negotiation

Although research on automated negotiation agents is still at a preliminary stage, extensive research on negotiation has been done using both game theory and behavioral approaches. Our discussion below focuses on how previous work in these two fields has been incorporated into electronic negotiation.

2.2.1 Game Theory Approach

Negotiation was treated as a bargaining game in early 1950s. The most famous one is two-person game, which can be classified into three categories: (1) pure conflict game, (2) pure coordination game, and (3) mixed-motive game (Tedeschi *et al.*, 1973). A pure conflict game exists when two parties have negatively correlated payoffs, while a pure coordination game exists when two parties have positively correlated payoffs. In a mixed-motive game, both conflict and coordination situations are present. This is similar to the concept of distributive versus integrative bargaining. Negotiation is characterized as a non-zero sum two-person game.

An aspect of game theory, which has gained strong empirical support and which has been widely applied in electronic negotiation, is the Nash solution (Nash, 1950; Nash 1953). Its significance in negotiation stems from the observation that it can be used to predict the outcome of a negotiation (Neslin & Greenhalgh, 1983; Eliashberg *et al.*, 1986) and to test the fairness of a negotiation outcome. In many studies on electronic negotiation, the Nash solution has been used as a performance measure to validate the fairness of negotiation agreements (Zeng & Sycara, 1998; Goh *et al.*, 2000)

Although a game theory approach seems to be a feasible method to study negotiation, it has several limitations. First of all, the game theory approach often focuses on outcome rather than process (Lim & Benbasat, 1993; Goh *et al.*, 2000; Zeng & Sycara, 1998). Second, game theorists are interested in how a player should behave to be smarter, while researchers are interested in how a player actually behave during the process (Luce & Raiffa, 1957; Raiffa, 1982). Thus, game-theory models are inadequate for research into understanding a negotiator's behavior during a negotiation process. Third, the number of players and their identities are fixed and known to everyone in game theory model. Also, all players are assumed to be rational, and their utility functions are disclosed (Zeng & Sycara, 1998). These assumptions might limit game theory's applicability to behaviorally-based negotiation.

2.2.2 Behaviorally-Driven Approach

To have autonomous negotiation, a negotiation model becomes a central component in electronic negotiation. A negotiation model is a set of rules or functions that agents follow to evaluate offers and make counter offers. A negotiation model that has been frequently employed is the time-dependent-tactical model, by which agents determine the values of the negotiating issues based on the time that has elapsed since the beginning of the negotiation. An example of negotiation agent employing this model is MIT's Kasbah (Kasbah, 1999). Here, agents negotiate with one another over the price of a product using one of the three time-decay functions: linear, quadratic or exponential (Guttman *et al.*, 1998).

However, when a negotiation involves multiple issues, it might become integrative because users might interest in different issues. To capture users' preferences on different negotiating issues, agents require utility functions to determine how much they like each received offer (Guttman *et al.*, 1998) and then to make counter offers by choosing the strategies that maximize the utilities. One utility function suggested by Sim & Chan (2000) is a linear combination formulation of the multi attribute utility theory (MAUT), which consists of the user-specified weight of each negotiating issue and the utility of each negotiating issue. The weights represent the importance levels of the issues to the users, and the utilities stand for the scores of the issues for the values in question. If a user has a higher concern on price than on delivery time, a higher weight is placed on price than delivery time. Another function suggested by Choi *et al.* (2001) differs from MAUT, in that the user-specified weight is the power term of the utility of each negotiating issue, and the overall utility is the cumulative product rather than the cumulative sum of the utilities of the negotiating issues. Among the studies we surveyed, the use of AHP for utility computation is, however, still absent. Since a negotiation process often involves multiple issues and multiple alternatives, and since AHP can be applied in selecting among competing alternatives in a multi-objective environment (Forman and Gass, 2001), we integrate AHP into the utility computation of negotiation agents.

Another gap in the behavioral approach to electronic negotiation is a lack of psychological elements. In human negotiation, two psychological factors that have been frequently studied are situational power and goal constraints. Psychologists have suggested that these two factors can affect a human's negotiation behavior (Pruitt, 1981; Rubin & Brown, 1975), and have tested their impacts on human negotiation outcome (Dwyer & Walker, 1981; McAlister *et al.*, 1986; Pinkley *et al.*, 1994; Huber & Neale, 1986; Bazerman *et al.*, 1985; Hamner & Harnett, 1974). However, research on incorporating situational power and goal constraints into negotiation agents is still at a preliminary stage.

2.2.3 Other Approaches

Besides game theory and behaviorally-driven approaches, learning an opponent's negotiation behavior has become another research direction in electronic negotiation. The forecast of an opponent's future concession might be used as a guideline for deciding a negotiator's upcoming move (Pruitt, 1981; Raiffa, 1982). This phenomenon is called tracking, and one way to implement it is through learning. Zeng & Sycara (1998) indicated that learning is beneficial for negotiation agents, because agents can achieve higher profits or utilities. Recently, two learning algorithms have been proposed for negotiation agents: (1) genetic algorithms, and (2) Bayesian learning.

A genetic algorithm is an evolutionary approach, in which high-performance negotiation tactics from parent agents are adapted to child agents, while low-performance tactics are removed. In the beginning, parent agents start with a population of negotiation tactics, and employ these tactics against other parent agents in negotiations. At the end of a negotiation, the performance of each tactic is evaluated. After evaluation, the tactics with high performance are integrated with other high-performance tactics, and are adapted to the child agents (Deveaux *et al.*, 2001;

Tu *et al.*, 2000). These crossed-over tactics of child agents are expected to outperform their parents'.

In Bayesian learning, each agent has some knowledge and belief about several aspects of a negotiation. These include the environmental or market situation, belief of opponents' negotiation tactics, and belief of opponents' utility functions. Moreover, before a negotiation starts, each agent has a subjective probability of each of this belief or knowledge. As the negotiation proceeds, these subjective probabilities are adjusted according to the received offers. With these updated subjective probabilities, agents make counter offers by following Bayesian decision-making rules (Zeng & Sycara, 1998).

Both genetic algorithms and Bayesian learning have limitations. To employ genetic algorithms, several trials have to be conducted before obtaining the child generation tactics (Deveaux *et al.*, 2001). In Bayesian learning, agents need to have prior knowledge before a negotiation starts. Thus, when agents are new to each other in a negotiation, both learning approaches may not be successful.

2.3 Our Contribution

In this thesis, our goal is to make negotiation agents to be more human-like, so that they can represent their owners in negotiations. To aid in this goal, we integrate learning and two psychological factors, namely situational power and goal constraints, into negotiation agents, and conduct experiments to test whether there is congruence of outcome between human and electronic negotiations.

In learning, we have observed that genetic algorithms and Bayesian learning have limited applicability when agents are new to each other in a negotiation. This is very important because there are numerous negotiation agents in the Internet, and agents might not have negotiation experience with each of their opponents. As a result, with incomplete information about an opponent, it is crucial to use the received offers in an established negotiation as a guideline for learning. Furthermore, among our reviewed studies, a learning algorithm for the time-dependent-tactical model is not available. Therefore, in this thesis, we introduce how we can use an agent's initial demand levels and concession rate for learning the time-dependent-tactical model's parameters.

As for situational power, we first find that the integration of this factor into electronic negotiation is missing. Second, our manipulation of situational power is also different from those in previous psychological studies. Instead of varying a negotiator's BATNA (Best Alternative to The Negotiated Agreement) (Pinkley *et al.*, 1994), or varying the number of buyers and sellers in the market (McAlister *et al.*, 1986; Dwyer & Walker, 1981), we use the idea of AHP's applicability in benchmarking an organization's performance (Forman & Gass, 2001) to model situational power. We make comparisons between the perceived market preferences of a product that buyers are interested in and a neutral preference, and use the difference to represent the level of power. Similarly, the incorporation of goal constraints into negotiation agents is absent, and our manipulation of goal constraints differs from those in previous psychological studies. Rather than telling a negotiator not to accept any transactions that do not meet the minimum profit requirements, we allow users to specify threshold utility values that agents must achieve before concluding the negotiations.

Finally, in the studies that we surveyed, we observe that the effects of situational power and goal constraints on electronic negotiation have not been tested. To address this, hypotheses tested in human negotiation are taken up for validation in the electronic setting. The goal here is not to propose new theories, but to test the level of agreement between human and electronic negotiations.

3 Time-Dependent Negotiation Tactic

In chapter 2, we reviewed numerous studies on electronic negotiation, and the gaps associated with it. We have also presented our contributions in the context of those studies. In the past decade, we observe that many different electronic negotiation agents such as MIT Media Lab's Kasbah (Kasbah, 1999), and MIT Media Lab's Impulse (Impulse, 1999), have been proposed and applied in electronic-commerce business.

With so many negotiation agents, Deveaux *et al.* (2001) have classified negotiation agents as either optimizing agents or behavioral negotiation agents. The objective of optimizing agents is to maximize their negotiation profits based on their beliefs of the market structure, and the expected behavior of their opponents. Thus, they do not follow certain preset models to negotiate. On the other hand, behavioral negotiation agents have an element called the type of behavior. With differences in the types of behavior, some agents like to make small changes in their initial offers, but concede quickly when approaching their time limits. In contrast, a few agents prefer to concede quickly at the beginning and stay firm at the end. In another case, some agents just keep changing their offers by a fixed amount along the negotiations. To capture these behaviors, behavioral negotiation agents employ preset negotiation protocols to decide how to reach concession or no-concession. Based on their negotiation models, behavioral negotiation agents follow their tactics to negotiate, until they reach their time limits or reservation values (the negotiating issues' thresholds at which negotiators would not go beyond).

In this thesis, we focus on behavioral negotiation agents, and how they can employ their negotiation models to negotiate. A negotiation model is a set of rules that negotiation agents follow to make offers or counter-offers. A negotiation tactic, which is a set of functions for agents to determine the values of the negotiating issues based on a single criterion (Matos *et al.*, 1998), is a major component of the negotiation model.

In the past few years, many different negotiation tactics have been proposed. Faratin *et al.* (1997) have identified the resource-dependent tactic, which generates offers based on the consumption pattern of resources during negotiation. When there are fewer resources, agents are more eager to complete the negotiations. On the other hand, the more the resources, the lower the pressure, and hence, agents are less urgent to reach agreements. Axelrod (1984) has introduced the behavior-dependent tactic, which generates offer according to opponents' previous attitudes. Negotiators simply imitate the behavior of their opponents for their next moves. Zacharia *et al.*, (2001) have also proposed a dynamic pricing algorithm, in which agents update the reputation of their counterparts in a collaborative fashion, and make offers based on these reputation values. Among these negotiation tactics, the one that has been frequently studied is the time-dependent tactic (Guttman *et al.*, 1998; Deveaux *et al.*, 2001; Faratin *et al.*, 1997; Matos *et al.*, 1998). With this tactic, agents treat time as the predominant factor, and determine the values of the negotiating issues in line with time pressure. In this chapter, we discuss the characteristics and the importance of this tactic.

3.1 Framework of Time-Dependent Tactic

The amount of time that has elapsed since the beginning of a negotiation creates time pressure, which forces negotiators to reach agreements quickly. Two effects of time pressure are lower demands and faster concessions (Pruitt, 1981). Since negotiators want to end the negotiations quickly, they have to sacrifice by accepting deals that are more favorable to their opponents. Rubin and Brown (1975) also stated that time pressure can cause greater or more frequent concessions. On the other hand, as negotiators are willing to wait longer and to appear less eager for reaching agreements, they achieve higher profits (Raiffa, 1982).

Pruitt (1981) and Raiffa (1982) have identified two types of time-dependent negotiation behavior, namely Boulware and Conceder. A negotiator with boulware behavior stays firm at the beginning of a negotiation, and concedes when approaching the time limit. When a negotiator starts with firmness, there is sufficient room to move, and the reservation value can be protected. In addition, a negotiator with boulware behavior can avoid position loss, which is the desertion of a desirable alternative, and also avoid image loss, which is a lack of firmness developed in other people's eyes. Position loss is a concern, because a negotiator has difficulty to withdraw a concession once made. Image loss is another concern, because once a negotiator concedes, it might inspire the opponent to maintain high demand. Negotiators with weak needs for agreements often exhibit boulware behavior, because they have less to lose by failure to reach agreements.

A negotiator with conceder behavior concedes his/her reservation values quickly at the beginning of a negotiation, and their concession rates become flattened as approaching the time limits. One reason for this behavior is to encourage the other parties to remain in the negotiation process so that agreements are more likely reached. Negotiators with strong needs for agreements usually exhibit conceder behavior, because they have more to lose if they fail to make agreements.

Based on the above concepts, Faratin *et al.* (1997) and Deveaux *et al.* (2001) have identified the mathematical expressions of the time-dependent tactic as follows:

$$P(t) = P_{\min} + e^{\left(1-\frac{t}{T_{\max}}\right)^{\beta} \ln(K)} \left(P_{\max} - P_{\min}\right) \qquad \text{If } V \text{ is decreasing} \qquad (3.1)$$

$$P(t) = P_{\min} + \left(1 - e^{\left(1-\frac{t}{T_{\max}}\right)^{\beta} \ln(K)}\right) \left(P_{\max} - P_{\min}\right) \qquad \text{If } V \text{ is increasing} \qquad (3.2)$$

where t is time (number of turn) in the interval [0, $T_{\max}$]

$P(t)$ is the value of the negotiating issue proposed by an agent at time t

$P_{\min}$ is the minimum value of the negotiating issue

$P_{\max}$ is the maximum value of the negotiating issue

K is a constant that determines the value of the negotiating issue in the first offer. $K \in (0, 1)$

$T_{\max}$ is the time limit (maximum number of turn) proposed by an agent

V is the utility function of the negotiating issue

β is a constant that determines the degree of convexity

To simplify our further analysis, we assume that an increase in the value of an issue decreases a buyer's utility, but increases a seller's utility. Under this assumption, equations (3.1) and (3.2) represent the time-dependent tactic expressions for a buying and a selling agent, respectively. In either equation, there are five major parameters, P_{min}, P_{max}, T_{max}, β, and K, specified by the agent's user. The interval [P_{min}, P_{max}] represents the value range of the issue that is acceptable by the user. A deadline, which is given by T_{max} in (3.1) and (3.2), represents the user's preferred time duration. The value of the issue in the first offer is obtained by multiplying the constant, K, by the interval [P_{min}, P_{max}]. The higher the value of K, the further the value of the issue in the first offer from the starting value. In addition, $K \in (0, 1)$ so that the value of the issue is always within the interval [P_{min}, P_{max}].

The parameter, β, is the degree of convexity that determines the type of behavior of the time-dependent tactic. The next section discusses how β can be changed to simulate different behavior.

3.2 Importance of Time-Dependent Tactic

We focus on time-dependent tactic in this thesis because of its similarity to human behavior, and its applicability in electronic negotiation.

3.2.1 Similarity to Human's Negotiation Behavior

As discussed in §3.1, time often has an impact on a negotiator's concession level. The mathematical model of the time-dependent tactic introduced in §3.1 can exhibit the effect of time pressure on a human's negotiation behavior.

Figure 3.1 shows the mathematical expression (3.1) with different values of β. It also shows the effect of β on the negotiation behavior of a buyer, when price is treated as the negotiating issue. Boulware behavior is exhibited when $\beta < 1$. The concession rate remains low until time is almost exhausted, where the value of the issue is conceded up to the reservation value. This is consistent with the thoughts from Raiffa (1982) and Pruitt (1981). The buyer makes lower initial demands, but faster concessions when approaching the time limit.

When $\beta > 1$, we observe conceder behavior from Figure 3.1. The concession rate is high at the beginning, but becomes low as the deadline approaches. This is similar to Pruitt's unilateral concession model, which is to concede unilaterally to reduce the distance between two parties' demands.

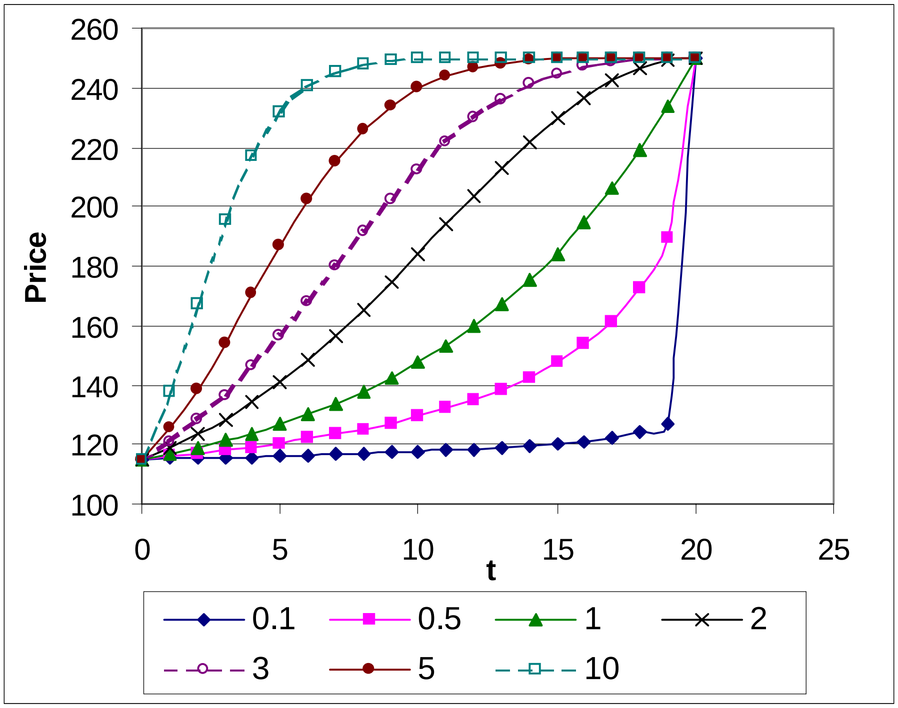

Figure 3.1. Effect of β on Time-Dependent Tactic

3.2.2 Applicability in Electronic Negotiation

Due to its resemblance to the human's negotiation behavior, the time-dependent tactic has been employed in many electronic negotiation agents such as Kasbah from MIT Media Lab (Kasbah, 1999; Guttman *et al.*, 1998). Kasbah provides buyers with one of three negotiation behaviors: anxious, coolheaded, and frugal – corresponding to a linear, quadratic or exponential function, respectively for increasing the price offers over time.

Moreover, numerous studies in electronic negotiation were conducted by employing the time-dependent tactic. Deveaux *et al.* (2001) used the time-dependent tactic in their experiments to test if an agent can achieve better performance by adapting its negotiation tactic to the behavior of its opponent. They employed both the conceder and the boulware behaviors in their analysis.

Faratin *et al.* (1997) also conducted several experiments to empirically test the applicability of the time-dependent tactic by allowing agents with boulware ($\beta < 1$), conceder ($\beta > 1$) and linear ($\beta = 1$) behaviors to negotiate with one another. They found that as there was plenty of time (high T_{max}) or only short time (low T_{max}) for agents to negotiate, agents with linear behavior outperformed those with conceder and boulware behaviors. This is analogous to the thoughts from Pruitt (1981) and Raiffa (1982). As Raiffa (1982) stated, "*The party that negotiates in haste is often at a disadvantage*". Negotiators with conceder behavior usually reach agreements quickly, but they have to sacrifice by accepting deals that are more favorable to their opponents (Pruitt, 1981). On the other hand, agents with boulware behavior usually fail to reach agreements and thus perform poorly (Raiffa, 1982; Pruitt, 1981). For agents with linear behavior, they often make moderate demands and reach agreements at good levels of profit.

In this chapter, we have reviewed the concept and the significance of the time-dependent tactic. With a correspondence to the human's negotiation behavior and applicability to electronic negotiation, more agents are expected to employ the time-dependent tactic in the future. Thus, if an agent can learn and understand the time-dependent tactic of its opponents, performance might be improved. Therefore, we will focus on the learning aspect of the time-dependent tactic in the next chapter.

4 Heuristics for Learning

In Chapter 3, we have investigated how a behavioral negotiation agent can employ the time-dependent tactic to negotiate. In this chapter, we attempt to design a learning algorithm for the parameters of the time-dependent tactic.

Learning capability is becoming an important component of a negotiation model because it can increase the performance a negotiation agent (Zeng & Sycara, 1998). Learning, which is defined as the ability to detect an opponent's negotiation model, might help a negotiation agent to improve over time. The decision of how much to demand or concede often depends on the expectations about the other's ultimate demand and concession rate (Pruitt, 1981; Raiffa, 1982), and the results from learning could help to make up these expectations. However, learning is difficult because of the scarcity of information. Since we assume that agents are new to each other, the only available information to estimate an opponent's negotiation model is the bid offers from the opponent in the established negotiation. A lack of opponent's negotiation history creates difficulty in learning. This is the reason why genetic algorithms have limited applicability when agents lack negotiation experience of their opponents (Deveaux *et al.*, 2001).

Once the result from learning is obtained, an agent could adjust its negotiation tactic to gain advantage based on the prediction of an opponent's future move. The objective of this chapter is to present heuristic learning algorithms for the estimation of the parameters of the time-dependent-tactical model. The predictability of the learning mechanism will be tested in the next chapter.

First of all, the difficulty of estimating the parameters of a time-dependent-tactical model is discussed in §4.1. Because of this difficulty, we introduce the properties of the derivatives of the time-dependent-tactical model, and an estimation approach of the derivatives in §§4.2 and §4.3, respectively. Then in §4.4, we present heuristic enumeration algorithms that use the derivatives' properties and the derivative-estimation approach. Based on the result from learning, an agent might achieve higher performance by adjusting the parameters of its negotiation model. Thus, in §4.5, we present an algorithm for agents to react to the estimated parameters. Finally, in a negotiation, the number of offers that an agent receives for learning might vary. The effect of this on the proposed learning algorithms is addressed in §4.6.

Due to the similarity of the negotiation models between the buying and the selling agents, the discussion in this chapter will focus on buying agents, although it is also applicable, with appropriate changes, to selling agents as well.

4.1 Difficulty of Estimation

In this thesis, learning represents an agent's ability to estimate the parameters of an opponent's time-dependent-tactical model. The estimated parameters need to closely fit the price offers. Previous experience with the opponent is not taken into account. Thus, each agent is considered to have similar level of difficulty in each new negotiation. In order to estimate the five parameters of the time-dependent-tactical model, several offers from an opponent are required before learning can start. One obvious way to estimate these five parameters is regression, which is discussed in §4.1.1. Another way is to use the first few price offers to analytically find a unique set of parameters, and this approach is addressed in §4.1.2. Both approaches lead to equations that appear difficult to solve mathematically.

4.1.1 Regression

Since the time-dependent tactic equation (3.1) consists of an exponential term, we first observe that a straightforward application of linear regression would not be possible. Next, we try to transform the negotiation model into a linear function. Chatterjee & Price (1991) have introduced several transformation methods for this purpose, but equation (3.1) is not amenable to any of these transformations because of the P_{min} term. The last approach we consider is nonlinear regression (Ratkowsky, 1983). To explain this, we write equation (3.1) as

$$P_t = P_{\min} + e^{\left(1-\frac{t}{T_{\max}}\right)^{\beta}\ln(K)}\left(P_{\max} - P_{\min}\right) = f\left(t,\theta\right) + \varepsilon_t$$

where $\theta = (P_{min}, P_{max}, \beta, T_{max}, K)^{T}$ is the vector of parameters to be estimated, t is the number of turns, ε_t is the error term that is normally distributed with mean zero and unknown variance, and the sum of squares, $S(\theta)$, which is to be minimized, is given by

$$S(\theta) = \sum_t [P_t - f(t,\theta)]^2 = [P_t - f(t,\theta)]^T [P_t - f(t,\theta)]$$

After substituting $f(t, \theta)$ with a one-degree Taylor's series polynomial, $S(\theta)$ can be differentiated and set to zero. The resultant equation can then be used to iteratively estimate the parameters by the Gauss-Newton method. Since there are five unknowns, we take t = 1, 2, 3, 4 and 5. That is, θ_{i+1}, the values of parameter set at iteration i+1, is given by

$$\theta_{i+1} = \theta_i + \left[J^T(\theta_i)J(\theta_i)\right]^{-1} J^T(\theta_i)[P_t - f(\theta_i)]$$

where i represents the iteration starting from one, and the $J(\theta_i)$ is the 5 x 5 Jacobian matrix given below (the row in the matrix represents t).

$$J(\theta) = \begin{bmatrix} \frac{\partial f_1(\theta)}{\partial P_{\min}} & \frac{\partial f_1(\theta)}{\partial P_{\max}} & \frac{\partial f_1(\theta)}{\partial \beta} & \frac{\partial f_1(\theta)}{\partial T_{\max}} & \frac{\partial f_1(\theta)}{\partial K} \\ \frac{\partial f_2(\theta)}{\partial P_{\min}} & \frac{\partial f_2(\theta)}{\partial P_{\max}} & \frac{\partial f_2(\theta)}{\partial \beta} & \frac{\partial f_2(\theta)}{\partial T_{\max}} & \frac{\partial f_2(\theta)}{\partial K} \\ \frac{\partial f_3(\theta)}{\partial P_{\min}} & \frac{\partial f_3(\theta)}{\partial P_{\max}} & \frac{\partial f_3(\theta)}{\partial \beta} & \frac{\partial f_3(\theta)}{\partial T_{\max}} & \frac{\partial f_3(\theta)}{\partial K} \\ \frac{\partial f_4(\theta)}{\partial P_{\min}} & \frac{\partial f_4(\theta)}{\partial P_{\max}} & \frac{\partial f_4(\theta)}{\partial \beta} & \frac{\partial f_4(\theta)}{\partial T_{\max}} & \frac{\partial f_4(\theta)}{\partial K} \\ \frac{\partial f_5(\theta)}{\partial P_{\min}} & \frac{\partial f_5(\theta)}{\partial P_{\max}} & \frac{\partial f_5(\theta)}{\partial \beta} & \frac{\partial f_5(\theta)}{\partial T_{\max}} & \frac{\partial f_5(\theta)}{\partial K} \end{bmatrix}$$

Preliminary simulations suggest the limited applicability of nonlinear regression approach to our estimation problem. First, convergence cannot be guaranteed in all cases, and is dependent on the initial set of estimated parameters. When the initial set of estimated parameters is close to the actual values, convergence appears to increase. Moreover, since the nonlinear regression method involves the use of Taylor's series with only the first two terms, the accuracy of the estimated parameters is another concern, even if convergence is obtained. Thus, we conclude that the nonlinear regression approach appears to be worthwhile for conducting further tests, before it can be actually implemented.

4.1.2 Price-Equation Algorithm

With equation (3.1) and the five price offers from an opponent, it might be possible to analytically estimate the five parameters of the time-dependent-tactical model. For example, by using equation (3.1) at t = 0 and t = 1, and then re-arranging terms, we obtain

$$P_{\min} = \frac{\mathrm{P}_{\mathrm{t}=0} - KP_{\max}}{1-K} \quad , t = 0 \tag{4.1}$$

$$P_{\max} = \frac{P_{t=1} - P_{\min} + P_{\min} e^{\left(1-\frac{t}{T_{\max}}\right)^{\beta} \ln(K)}}{e^{\left(1-\frac{t}{T_{\max}}\right)^{\beta} \ln(K)}} \quad , t = 1 \tag{4.2}$$

When equations (4.1) and (4.2) are substituted into (3.1), we have

$$P(t) = \frac{P_{t=0} - \frac{K\left(P_{t=1}(K-1) + P_{t=0} - P_{t=0}E\right)}{K-E}}{1-K} + E\left(\frac{\left(P_{t=1}(K-1) + P_{t=0} - P_{t=0}E\right)}{K-E} - \frac{P_{t=0} - \frac{K\left(P_{t=1}(K-1) + P_{t=0} - P_{t=0}E\right)}{K-E}}{1-K}\right)$$

$$(4.3_t)$$

where $E = e^{\left(1-\frac{t}{T_{\max}}\right)^{\beta} \ln(K)}$.

Equation (4.3_t) consists of three equations for t = 2, 3, 4 and three unknowns, β, T_{max} and K. No further simplification could be done. We have tried to use the well known software, Maple (Kofler, 1997) to solve these equations symbolically but did not succeed. In order to solve three equations simultaneously, trial and error is required.

4.2 Properties of Derivative

Due to the limitations of regression and the price-equation algorithm, we look further into the properties of the time-dependent-tactical model. In this model, the dependent variable is price or the value of the negotiating issue, while two of the independent variables are the number of turns and the degree of convexity, β, which determines the behavior of the agent. The discussion below focuses on the properties of the derivatives of the price with respect to these two parameters. The use of these properties in our algorithms is addressed in §4.4.

4.2.1 Derivative with respect to the Number of Turns

The first, second and third derivatives of the price with respect to the number of turns are shown in (4.4), (4.5) and (4.6), respectively.

$$\frac{\partial P}{\partial t} = \frac{-(P_{max} - P_{min})\beta \ln(K)}{T_{max}} \left[\frac{e^{\left(1-\frac{t}{T_{max}}\right)^{\beta} \ln(K)} \left(1-\frac{t}{T_{max}}\right)^{\beta}}{\left(1-\frac{t}{T_{max}}\right)} \right] \tag{4.4}$$

$$\frac{\partial^2 P}{\partial t^2} = C_1 (C_2 + C_3), \tag{4.5}$$

where $C_1 = \frac{(\mathrm{P}_{max} - \mathrm{P}_{min})\beta \ln(K)}{\mathrm{T}_{max}{}^2} \left[\frac{\mathrm{e}^{\left(1-\frac{\mathrm{t}}{\mathrm{T}_{max}}\right)^{\beta} \ln(K)} \left(1-\frac{t}{T_{max}}\right)^{\beta}}{\left(1-\frac{t}{T_{max}}\right)^2} \right]$, $C_2 = \beta \ \ln(K) \left(1-\frac{t}{T_{max}}\right)^{\beta}$,

$C_3 = (\beta - 1)$.

$$\frac{\partial^3 P}{\partial t^3} = \mathrm{C}\left[-\beta^2 + 3\beta - 3\beta^2 \ln(K)\left(1-\frac{t}{T_{\max}}\right)^{\beta} - 2 + 3\beta \ln(K)\left(1-\frac{t}{T_{\max}}\right)^{\beta} - \beta^2 (\ln(K))^2 \left(1-\frac{t}{T_{\max}}\right)^{2\beta} \right] \tag{4.6}$$

where $\mathrm{C} = \dfrac{(\mathrm{P}_{\max} - \mathrm{P}_{\min})\beta \ln(K)\left(1-\frac{t}{T_{\max}}\right)^{\beta} e^{\left(1-\frac{t}{T_{\max}}\right)^{\beta}\ln K}}{T_{\max}{}^{3}\left(1-\frac{t}{T_{\max}}\right)^{3}}$.

Lemma 1. *For all $x > 0$, the value of $\left(1-\frac{t}{T_{\max}}\right)^{x}$ monotonically decreases from 1 to 0 as t increases from 0 to T_{max}.*

Proof Clearly, when $x > 0$, $\left(1-\frac{t}{T_{\max}}\right)^{x}$ is 1 and 0 for $t = 0$ and $\mathrm{T}_{\max}$, respectively. Further, $\left(1-\frac{t_1}{T_{\max}}\right) > \left(1-\frac{t_2}{T_{\max}}\right)$ for $t_1 < t_2$. This shows the Lemma.

Corollary 1. For $K \in (0, 1)$, C_2 is negative, and increases monotonically from $\beta\ln(K)$ to 0 as t increases from 0 to T_{max}.

Theorem 1. *The price of a buying agent increases with the number of turns.*

Proof Since the parameters (P_{min}, P_{max}, β, and T_{max}) are positive, and $\ln(K)$ is negative for $K \in (0, 1)$, it follows that the term outside the bracket in equation (4.4) is positive. From *Lemma 1*, the term inside the bracket is greater than or equal to zero. Thus, this theorem is proved.

Lemma 2. *For $t \in [0, T_{max}]$, $C_1 \leq 0$.*

Proof Similar to that for *Theorem 1*.

Theorem 2. *For boulware agents, the first derivative of price with respect to the number of turns is always increasing.*

Proof Because $C_1 \leq 0$ by *Lemma 2*, $C_2 \leq 0$ by *Corollary 1*, and $C_3 \leq 0$ for boulware agents, $\frac{\partial^2 P}{\partial t^2} \geq 0$.

Theorem 3. *For conceder agents, there exists a $t \in [0, T_{max}]$ after which $\frac{\partial P}{\partial t}$ decreases.*

Proof $C_1 \leq 0$ by *Lemma 2*, and $C_3 > 0$ for conceder agents. Hence, if for $t = 0$, $|C_2| = |\beta \ln(K)| < C_3$, then $\frac{\partial^2 P}{\partial t^2} < 0$. If not, because $C_2 \leq 0$ and monotonically increases to 0 by *Corollary 1*, there exists a t such that $|C_2| < C_3$. After that t, since C_2 increases monotonically, $\frac{\partial^2 P}{\partial t^2}$ remains negative, and $\frac{\partial^2 P}{\partial t^2} < 0$ in turn.

4.2.2 Derivative with respect to Concession Rate

Equation (4.7) below shows the first derivative of price with respect to β.

$$\frac{\partial P}{\partial \beta} = B_1 B_2, \tag{4.7}$$

where $B_1 = \left(1 - \frac{t}{T_{\max}}\right)^{\beta} \ln\left(1 - \frac{t}{T_{\max}}\right)$, $B_2 = e^{\left(1-\frac{t}{T_{\max}}\right)^{\beta} \ln(K)} (P_{\max} - P_{\min}) \ln(K)$.

Theorem 4. *At every turn, the price increases with* β

.

Proof According to *Lemma 1*, the first and the second terms in B_1 above are positive and negative respectively, and hence, $B_1 \leq 0$. The parameters (P_{min}, P_{max}, β, and T_{max}) are positive, ln(K) is negative for $K \in (0, 1)$, and $B_2 < 0$ for all t and β. As a result, $\frac{\partial P}{\partial \beta} \geq 0$ for all t and β.

4.3 Estimation of Derivatives

The learning algorithms that will be proposed in §4.4 are based on the properties of the derivatives introduced in §4.2, and on the approach to estimate these derivatives at various turns. This section focuses on the latter aspect.

There are two approaches for the estimation. The first approach consists of taking the derivative between two turns by connecting the two price offers with a straight line. Since the time difference between consecutive turns is always one, the derivative could be estimated as the difference between the prices. The advantage of this method is the ease of calculation. The disadvantage is the accuracy of the estimated derivative. The time-dependent-tactical model is a curve, while the above approach relies on connecting two price offers with a straight line. This difference makes this estimation method inaccurate.

Our second approach applies the Taylor's series for the estimation. When the time-dependent-tactical model is approximated by a Taylor's series, the coefficients of the series represent the derivatives. The use of Taylor's series is similar to that in the nonlinear regression approach. Ratkowsky (1983) employs a first-degree Taylor's polynomial to describe a nonlinear regression method, whereas we use a third-degree Taylor's polynomial. Later, we will explain why we choose third-degree polynomials. With the mathematical tractability of the derivatives at $t = 0$, we choose to approximate the negotiation model, $P(t)$, at $t = 0$, and the Taylor's series polynomial is given as

$$\hat{P}(t) = P(0) + P'(0)t + \frac{P''(0)}{2}t^2 + \frac{P'''(0)}{3!}t^3 + \ldots + \frac{P^n(0)}{n!}t^n + R_{n+1}$$

where $\hat{P}(t)$ is the estimated price at t, $P(0)$ is the actual price at t = 0, n is the number of degree of the polynomial, $P'(0), P''(0), P'''(0), P^n(0)$ are the various derivatives of the price with respect to the number of turns at t = 0 and at various degree of differentiation, and $R_{n+1} = \frac{P^{n+1}(c)}{(n+1)!}t^{n+1}$ for $c \in [0, t]$.

The Taylor's series expression allows us to work with a polynomial approximation of the time-dependent-tactical model. Theoretically, before such an expression can be used, we need to determine the degree of the polynomial for which the remainder converges to zero in the limit. However, due to the complexity of the derivatives, we opt to employ simulation to determine the degree of the polynomial that must be used, rather than prove convergence mathematically.

We conducted a preliminary simulation to test four different degrees of polynomials (1, 2, 3 and 4). Let SSE_{1st} denotes the sum of the squared error (*SSE*) between the estimated first derivative obtained from the Taylor's polynomial and the actual value computed by (4.4), SSE_{2nd} means the *SSE* between the estimated second derivative from the polynomial and the actual one from (4.5), and SSE_{3rd} stands for the *SSE* between the estimated third derivative taken from the polynomial and the actual third derivative from (4.6). The three *SSEs* are presented in Table 4.1.

Table 4.1. The SSE_{1st}, SSE_{2nd}, SSE_{3rd} of Four Polynomial Models.

Performance Measure	Number of Degree of Polynomial			
	One	**Two**	**Three**	**Four**
SSE_{1st}	8.6072	1.1681	0.0024	0.0007
SSE_{2nd}	6.1124	2.4861	0.0217	0.0104
SSE_{3rd}	1.1135	1.1135	0.0617	0.0450
SSE_{1st}, excluding t = 0	4.1214	0.6081	0.0018	0.0004
SSE_{2nd}, excluding t = 0	3.7581	1.4541	0.0159	0.0038
SSE_{3rd}, excluding t = 0	0.9379	0.9379	0.0411	0.0144

According to Table 4.1, the third-degree and the fourth-degree polynomials outperform the other two in three *SSEs*. An example of Taylor's polynomial fits for the first five price offers of a buying agent (P_{min} = 100, P_{max} = 250, β = 5.5, T_{max} = 10, K = 0.1) is given in Figure 4.1. As shown therein, all polynomials fit the price offers well. However, according to the first derivative plot, the third-degree and the fourth-degree polynomials outperform the first-degree and the second-degree polynomials. Further, a t-test between the third-degree and the fourth-degree polynomials indicates that there are insignificant differences between these two polynomial approximations for all the three *SSEs*. Summers *et al.* (1977) stated that an appropriate rule was to select the lowest-order polynomial, when more than one polynomial with different degrees appeared to fit the series reasonably well. Hence, we select the third-degree polynomial. Based on our empirical experience, the third-degree polynomial fits well for a variety of different types of agents.

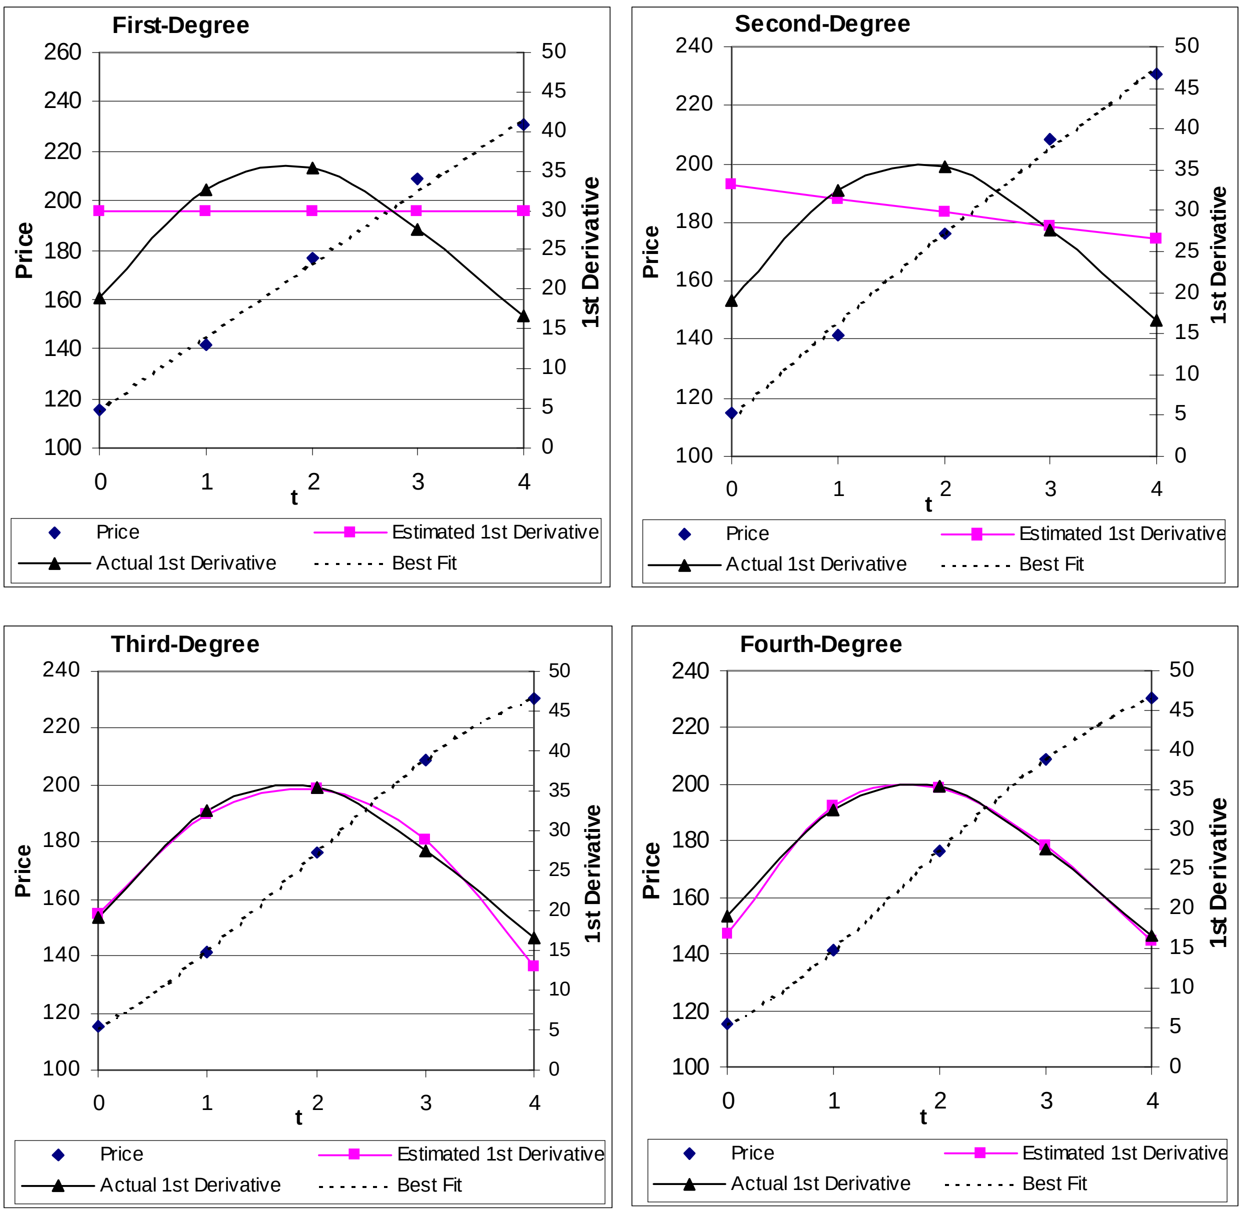

Figure 4.1. Samples of Taylor's Polynomial Fits for the First Five Price Offers.

For our simulation, Table 4.1 reveals a pattern that for the third-degree polynomial, SSE_{1st} is smaller than SSE_{2nd}, while SSE_{2nd} is smaller than SSE_{3rd}. That is, for higher-order derivatives, the accuracy of the estimated derivative decreases. In addition, when the error term at $t = 0$ is not taken into account in *SSE* computation, we find significant decreases in three *SSEs* of the third-degree polynomial. The estimated derivatives at $t = 0$ are comparatively less accurate than the estimated derivatives at other values of t. Finally, the simulation result is only used as a heuristic.

Based on the foregoing discussion, we approximate the time-dependent-tactical model by

$$\hat{P}(t) = b_0 + b_1 t + b_2 t^2 + b_3 t^3 \tag{4.8}$$

where $\hat{P}(t)$ is the estimated price at turn t, and b_0, b_1, b_2 and b_3 are the terms obtained from the best-fit curve. By differentiating (4.8), we obtain:

$$\hat{P}'(t) = b_1 + 2b_2 t + 3b_3 t^2 \tag{4.9}$$

$$\hat{P}''(t) = 2b_2 + 6b_3 t \tag{4.10}$$

$$\hat{P}'''(t) = 6b_3 \tag{4.11}$$

where $\hat{P}'(t), \hat{P}''(t),\ \hat{P}'''(t)$ are the estimated first, second and third derivatives at t.

4.4 Algorithms

As stated in §4.1, there is a difficulty to estimate the parameters by analytically solving the time-dependent tactic equations. This section introduces how we use the properties of the derivatives and the estimates of the derivatives to define an enumeration range. Our heuristic learning algorithms consist of three phases: 1) Preprocessing Phase, 2) Construction Phase, and 3) Improvement Phase. To varying degrees, these phases look for a balance between accuracy and tractability. The preprocessing phase seeks to reduce the search range of β by categorizing an opponent's behavior as either boulware or conceder. In the construction phase, a set of parameters is estimated based on the approximate values of $\frac{\partial P}{\partial t}$ at t = 0. Finally, the estimated parameters from the construction phase can then be used as a reference for defining the search range in the improvement phase. We then perform a discretized enumeration within that range to estimate the parameters.

4.4.1 Preprocessing Phase

The objective of the preprocessing phase is to classify whether an agent is conceder or boulware based on Theorems 2 and 3, and on estimating the time-dependent-tactical model using a third-degree best-fit curve. This classification enables us to specify the range of β in the construction phase.

The pseudo-code of the preprocessing phase is presented in Figure 4.2. First, a third-degree best-fit curve, with the format of equation (4.8), is constructed based on the price offers from the opponent (Step 1). After that, we compute the $\hat{P}'(t)$ and the $\hat{P}''(t)$ for t = 0 to the number of price offers minus one, N – 1 (Step 2). According to Theorem 2, if $\frac{\partial P}{\partial t}$ decreases, the agent cannot exhibit boulware characteristic. Furthermore, Theorem 3 states that $\frac{\partial P}{\partial t}$ must decrease for a conceder agent. As a result, when $\hat{P}'(t)$ decreases at any t, the agent is said to be conceder. This allows us to set the search range as given in Step 3.

Preprocessing Phase

1) Construct a cubic best-fit polynomial from the price offers with the format of equation (4.8)
2) Find the $\hat{P}'(t)$ and $\hat{P}''(t)$ from t = 0 to (N-1) by equations (4.9) and (4.10)
3) If $\hat{P}'(t+1) < \hat{P}'(t)$ for same t
    Set β_{lower} = 1.01, since agent is conceder
  Else
    Set β_{lower} = 0.01, since agent might be boulware or conceder
  End If
  **Note: The accuracy of β is set to 2 decimal places*

Figure 4.2. Pseudo-code of the Preprocessing Phase.

4.4.2 Construction Phase

The objective of the construction phase is to find out a set of estimated parameters for reference, based on the derivatives of the price at t = 0. Derivatives at t = 0 are used in the construction phase, because the resulting equations can be simplified. We propose two algorithms: 1) The Two-Derivative Algorithm and 2) The Three-Derivative Algorithm. The former one employs the first two derivatives at t = 0, while the latter one employs the first three derivatives at t = 0.

4.4.2.1 The Two-Derivative Algorithm

Here, we combine the first and the second derivatives at t = 0 to obtain

$$T_{\max} = \frac{-\left.\frac{\partial P}{\partial t}\right|_{t=0}}{\left.\frac{\partial^2 P}{\partial t^2}\right|_{t=0}}\left[\beta \ln K + \beta - 1\right]$$

This can be simplified further by using the derivatives' estimates, as given below.

$$T_{\max} = \frac{-\hat{P}'(0)}{\hat{P}''(0)}\left[\beta \ln K + \beta - 1\right] \qquad (4.12)$$

The resulting values of parameters can then be combined with (4.2) and (4.1) to evaluate P_{max} and P_{min}. In this algorithm, several sets of parameters are obtained by enumerating over a discretized range of β and K. The estimated set of parameter with the least sum of squared error (*SSE*) between the actual prices and the estimated prices is selected as the reference set.

The pseudo-code of this algorithm is presented in Figure 4.3. To perform the above computation, we enumerate over the feasible ranges of K and β (Steps 2a and 2b), and then find out T_{max}, P_{min} and P_{max} as given in Steps 2c and 2d. This enumeration is continued, until the *SSE* reaches a minimum and starts to increase. The stopping criteria stems from our empirical experience in which we often observed the *SSE* first decreased before increasing (See Figure 4.4). Furthermore, we observed that when *SSE* started to increase, the previous β was the appropriate value for the corresponding K. We then compare the *SSE* with the current minimum

SSE, and update the selected parameters. On the other hand, if *SSE* did not start to increase at $\beta = \beta_{limit}$, the solutions obtained were often far from the actual values. Hence, we do not use the solution (Steps 2e, 2f and 2g). This enumeration is continued, until all values of *K* are processed.

Construction Phase – The Two-Derivative Algorithm

1) Initialization

    Set $\hat{P}'(0)$ and $\hat{P}''(0)$ from preprocessing phase
    Set $K_{lower} = 0.1$ and $K_{upper} = 0.9$
    Set β_{lower} and β_{limit} from preprocessing phase

2) Construction

    (a) For $K = K_{lower}$ to K_{upper} with an incremental level of 0.1
    (b) For $\beta = \beta_{lower}$ to β_{limit} with an incremental level of 0.01
    (c) Find T_{max} by equation (4.12)
        If $T_{max} > N$-1
    (d) Estimate P_{max} and P_{min} by equations (4.2) and (4.1)
        If $P_{max} > 0$, $P_{min} > 0$, and $P_{max} > P_{min}$,
    (e) Find $P_{estimated}(t)$ from $t = 0$ to N-1

$$SSE(K,\beta) \equiv \sum_{t=0}^{N-1} \left(P(t) - P_{estimated}(t)\right)$$

    (f) If $SSE(K, \beta) < SSE(K, \beta\text{-}0.01)$ and $\beta \geq \beta_{limit}$
            Set $SSE(K)$ = no solution
        ElseIf $SSE(K, \beta) \geq SSE(K, \beta\text{-}0.01)$
            Set $SSE(K) = SSE(K, \beta\text{-}0.01)$
            Exit β loop
        End If
        End If
        End If
        Next β
    (g) Set $SSE_{\text{Two-Derivative}} = \min\{SSE(K)\}$
        Update the selected parameter set $(\beta, T_{max}, K, P_{min}, P_{max})_{\text{Two-Derivative}}$ corresponding to that minimum
    Next *K*

Figure 4.3. Pseudo-code of the Construction-Phase's Two-Derivative Algorithm.

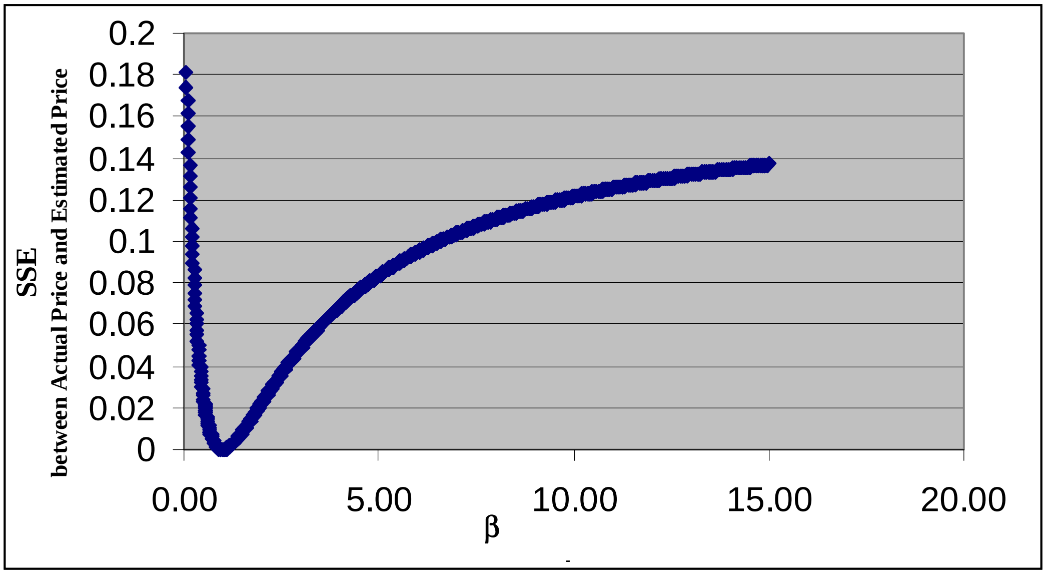

Figure 4.4. Effect of β on the *SSE* for the First Five Price Offers.

4.4.2.2 The Three-Derivative Algorithm

Here we first combine the first and the third derivatives at $t = 0$ to obtain

$$T_{\max}{}^2 = \frac{-\left.\frac{\partial P}{\partial t}\right|_{t=0}}{\left.\frac{\partial^3 P}{\partial t^3}\right|_{t=0}} \left[-\beta^2 + 3\beta - 3\beta^2 \ln(K) - 2 + 3\beta \ln(K) - \beta^2 (\ln(K))^2\right] \qquad (4.13)$$

Then, when we combine (4.12), which consists of the first and the second derivatives at $t = 0$, and (4.13), we obtain a quadratic equation of β.

$$M_1\beta^2 + M_2\beta + M_3 = 0$$

where $M_1 = (a-b)(\ln(K))^2 + (a-b) + (2a-3b)\ln(K)$, $M_2 = (3b-2a)\ln(K) + (3b-2a)$,

$M_3 = a - 2b$, $a = \dfrac{\left.\frac{\partial P}{\partial t}\right|_{t=0}}{\left(\left.\frac{\partial^2 P}{\partial t^2}\right|_{t=0}\right)^2}$, and $b = \dfrac{1}{\left.\frac{\partial^3 P}{\partial t^3}\right|_{t=0}}$.

The values of a and b can be obtained by the derivatives' estimates, as given in (4.14). This yields M_1, M_2 and M_3, and in turn the value of β from the quadratic equation.

$$a = \frac{\hat{P}'(0)}{(\hat{P}''(0))^2},$$
$$b = \frac{1}{\hat{P}'''(0)} \tag{4.14}$$

As with the two-derivative algorithm, the resulting values of parameters can then be combined with (4.2) and (4.1) to evaluate P_{max} and P_{min}. In this algorithm, several sets of parameters are obtained by enumerating through a defined range of K. The estimated set of parameters with the least sum of squared error between the actual prices and the estimated prices is selected as the reference set.

The pseudo-code of this algorithm is presented in Figure 4.5. To perform the above computation, we enumerate over the feasible range of K (Step 2a), and then find out the coefficients as well as the determinant, Δ, as given in Steps 2b and 2c. If $\Delta < 0$, no solution is possible (Step 2d). If $\Delta \geq 0$, β is found by using the quadratic formula (Step 2e). If β is positive, we evaluate T_{max}, P_{max}, P_{min} and corresponding *SSE* as given in Steps 2f, 2g and 2h. After that, we compare the *SSE* with the current minimum *SSE*, and update the selected parameters (Step 2i).

Construction Phase – The Three-Derivative Algorithm

1) Initialization

Set $\hat{P}'(0)$ and $\hat{P}''(0)$ from preprocessing phase

Set $K_{lower} = 0.1$ and $K_{upper} = 0.9$

Find $\hat{P}'''(0)$ by equation (4.11)

2) Construction

(a) For $K = K_{lower}$ to K_{upper} with an incremental level of 0.1

(b) Find the coefficients (M_1, M_2 and M_3) of the quadratic equation (4.14)

(c) Find the determinant, $\Delta = (M_2)^2 - 4M_1M_3$

If $\Delta < 0$

(d) Set $SSE(K)$ = no solution

ElseIf $\Delta \geq 0$

(e) Find β_i by using quadratic formula (The subscript i is used to differentiate two β if $\Delta > 0$)

For each β_i

If $\beta_i > 0$

(f) Find T_{max} by using equation (4.12)

If $T_{max} > N-1$

(g) Find P_{max} and P_{min} by equations (4.2) and (4.1)

If $P_{max} > 0$, $P_{min} > 0$, and $P_{max} > P_{min}$

Find $P_{estimated}(t)$ from t = 0 to N-1

(h) $SSE(K) \equiv \sum_{t=0}^{N-1} (P(t) - P_{estimated}(t))$

End If

End If

End If

Next β_i

(i) Set $SSE_{\text{Three-Derivative}} = \min\{SSE(K)\}$

Update the selected parameter set $(\beta, T_{max}, K, P_{min}, P_{max})_{\text{Three-Derivative}}$ corresponding to that minimum

End If

Next K

Figure 4.5. Pseudo-code of the Construction-Phase's Three-Derivative Algorithm.

4.4.3 Improvement Phase

In §4.1.2, we saw that the difficulty in estimating the parameters of the time-dependent-tactical model is due to the lack of a limited search range. This problem can be mitigated by using the set of estimated parameters from the construction phase as a reference, and then defining a discretized range of parameters to search.

We propose three algorithms: 1) The Price Algorithm, 2) The First-Derivative Algorithm, and 3) The Second-Derivative Algorithm. These algorithms have a similar structure, but differ in the criteria used to guide the estimation. The price algorithm employs the prices, while the first-derivative and the second-derivative algorithms employ the estimated first derivatives and the estimated second derivatives from the best-fit-third-degree polynomial, respectively.

Based on the reference set of parameters $\left(K^{ref}, \beta^{ref}, T_{\max}^{ref}\right)$ from the construction phase, we define the search ranges of β, T_{max} and K as presented in Tables 4.2, 4.3 and 4.4 There is no standard to define the search range based on a reference value. The greater the search range, the more accurate the estimated parameters, but the longer the searching time. In this thesis, we define the search range for β according to the observed behavior type, and T_{max} based on short, medium, and long term deadlines.

Table 4.2. Search Range for β in the Improvement Phase.

β^{ref}	Observed Behavior	β_{lower}	β_{upper}
$\beta^{ref} \leq 0.5$	Highly Boulware	0.01	1.00
$0.5 < \beta^{ref} < 1$	Moderately Boulware	0.10	$2\ \beta^{ref}$
$1 \leq \beta^{ref} \leq 3$	Linear/Slightly Conceder	1.00	$1.5\ \beta^{ref}$
$3 < \beta^{ref} \leq 5$	Moderately Conceder	$0.5\ \beta^{ref}$	$1.5\ \beta^{ref}$
$\beta^{ref} > 5$	Highly Conceder	$0.75\ \beta^{ref}$	$1.25\ \beta^{ref}$

Note: The incremental level of β is always 0.01 in this thesis.

Table 4.3. Search Range for T_{max} in the Improvement Phase.

$T_{\max}^{ref}$	Characteristic	$T_{max\ lower}$	$T_{max\ upper}$
$T_{\max}^{ref} \leq 10$	Short Term Deadline	5	$2\ T_{\max}^{ref}$
$10 < T_{\max}^{ref} \leq 30$	Medium Term Deadline	$0.5\ T_{\max}^{ref}$	$1.5\ T_{\max}^{ref}$
$T_{\max}^{ref} > 30$	Long Term Deadline	$0.5\ T_{\max}^{ref}$	$1.25\ T_{\max}^{ref}$

Note: The incremental level of T_{max} is always one in this thesis.

Table 4.4. Search Range for K in the Improvement Phase.

K^{ref}	K_{lower}	K_{upper}
$0 < K^{ref} < 1$	$K^{ref} - 0.2$	$K^{ref} + 0.2$

Note: The incremental level of K is always 0.1 in this thesis.
When $K_{lower} < 0.1$, $K_{lower} = 0.1$
When $K_{upper} > 0.9$, $K_{upper} = 0.9$

The pseudo-code of the improvement phase is presented in Figure 4.6. With the search ranges, we next enumerate over K, β, and T_{max} (Step 1), and find out P_{max} and P_{min} as given in Step 2. After that, we compute the *SSE* (see below for more explanation), and choose the parameters with the lowest *SSE* (Step 3).

For the price algorithm, we compute the estimated price, $P_{estimated}(t)$, at t = 2 to N-1, by plugging into (3.1) the estimates of P_{max} and P_{min} (Step 3b of Figure 4.6a) and the values of K, β and T_{max} (Step 2 of Figure 4.6a). With these estimates, we compute the SSE between the actual price offers and the $P_{estimated}(t)$ as given in Step 2 of Figure 4.6b. The first- and the second-derivative algorithms are similar in idea, except that the values of K, β and T_{max} and the estimates of P_{max} and P_{min} are substituted into (4.4) and (4.5), respectively, as shown in Figure 4.6b.

Improvement Phase

1) Initialization

Set $\left(K^{ref}, \beta^{ref}, T_{\max}^{ref}\right)$ = estimated parameter set from the construction phase

2) Defining

Set K_{lower}, and K_{upper} based on K^{ref} and Table 4.2

Set β_{lower}, β_{upper} based on β^{ref} and Table 4.3

Set $T_{max\ lower}$, $T_{max\ upper}$, based on $T_{\max}^{ref}$ and Table 4.4

3) Estimating

```
(a) For K = K_lower to K_upper with an incremental level of 0.1
      For T_max = T_max lower to T_max upper with an incremental level of 1
        For β = β_lower to β_upper, with an incremental level of 0.01
(b)         Find P_max and P_min by equations (4.2) and (4.1)
            If P_max > 0, P_min > 0, and P_max > P_min
(c)            Compute SSE as in Figure 4.6b
(d)            If SSE(K, T_max, β) < minimum
                  Set minimum = SSE(K, T_max, β)
                  Set (P̂_min, P̂_max, T̂_max, β̂, K̂) = (β, T_max, K, P_min, P_max)
               End If
            End If
        Next β
      Next T_max
   Next K
```

Figure 4.6a. Pseudo-code for the Improvement-Phase's Algorithms.

The Price Algorithm	The First-Derivative Algorithm	The Second-Derivative Algorithm
1. Find the $P_{estimated}(t)$ from t = 2 to (N – 1) by (3.1)	1. Find the $\frac{\partial \hat{P}}{\partial t}$ from t = 1 to (N – 1) by (4.4)	1. Find the $\frac{\partial^2 \hat{P}}{\partial t^2}$ from t = 1 to (N – 1) by (4.5)
2. $SSE(K, T_{\max}, \beta) \equiv \sum_{t=2}^{N-1}\left(P(t) - P_{estimated}(t)\right)$	2. $SSE(K, T_{\max}, \beta) \equiv \sum_{t=1}^{N-1}\left(\hat{P}'(t) - \frac{\partial \hat{P}}{\partial t}\right)$	2. $SSE(K, T_{\max}, \beta) \equiv \sum_{t=1}^{N-1}\left(\hat{P}''(t) - \frac{\partial^2 \hat{P}}{\partial t^2}\right)$

Figure 4.6b. Pseudo-code for the SSE computation in the Improvement Phase.

Figure 4.6. Pseudo-code of the Improvement Phase.

4.5 Using the Estimated Parameters

The results from learning can be used to adjust the parameters of the time-dependent-tactical model. This is important because the estimated parameters can be used as a guideline for deciding an agent's concession level in the future. However, depending on the accuracy of learning, the adjustment might cause an agent to achieve higher or lower negotiation performance. This section introduces a *reaction* algorithm we use to change the values of the parameters of an agent's time-dependent-tactical model. This algorithm focuses on buying agents, but the essential logic is applicable, with suitable modification, to selling agents as well.

Our algorithm consists of three phases: (1) Selection of Target Range, (2) Feasibility Check, and (3) Adjustment. The purpose of the selection of target range is to find a range of estimated prices that might be beneficial to the agents, based on the estimated parameters. In the feasibility-check phase, we find out the best-estimated price that is attainable according to the parameters of an agent's current negotiation model. Finally, in the adjustment phase, we change the values of the parameters of an agent's negotiation model so that the agreement can be made at the desired estimated price. We next explain these three phases. The pseudo-code of our *reaction* algorithm is presented in Figure 4.11.

4.5.1 Selection of Target Range

The first step of our algorithm is to decide if we can achieve higher performance, based on the estimated parameters. With the estimated price offers and our negotiation model, we can predict the final price of the agreement if one is possible. Thus, by comparing this final price with the other estimated price offers, an agent can decide if a better deal is possible (Step 1c of Figure 4.11). Figure 4.7 shows a negotiation process involving a possible final agreement. The estimated parameters indicate that an agreement will be made at point B. If the agreement is made at other estimated prices between A and B, the buying agent can achieve higher performance.

In another situation, wherein an agreement is not possible based on the estimated parameters, we include all estimated prices for consideration. Figure 4.8 shows this situation. According to the estimated prices, both parties will not compromise for an agreement (i.e., t_i in Step 1d is not present). However, if the agreement is made at any estimated price between C and D, the buying agent can achieve higher performance. Hence, t_{end} is set to the turn corresponding to the highest-estimated offer, as given in Step 1f.

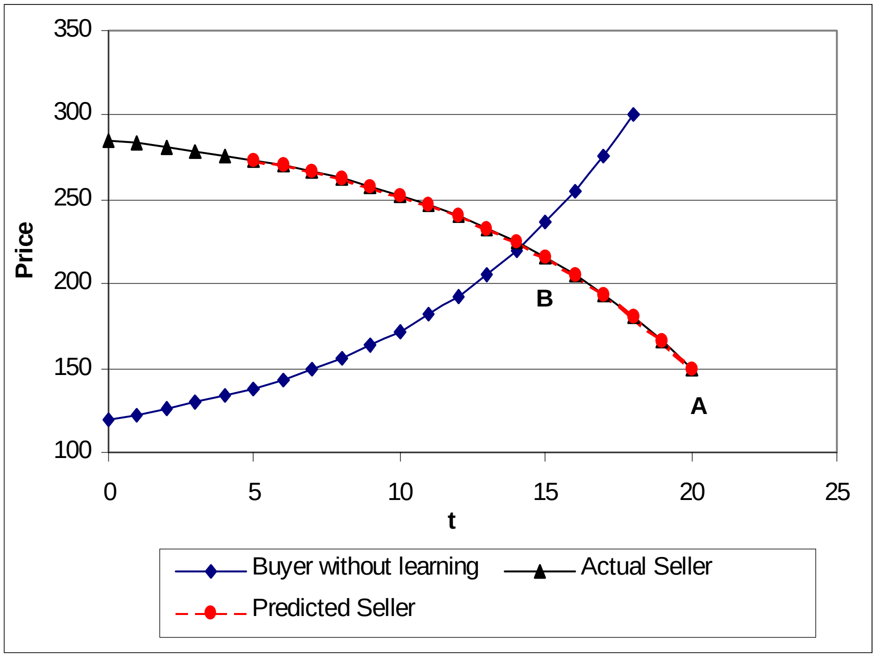

Figure 4.7. A Negotiation Process involving a Possible Final Agreement.

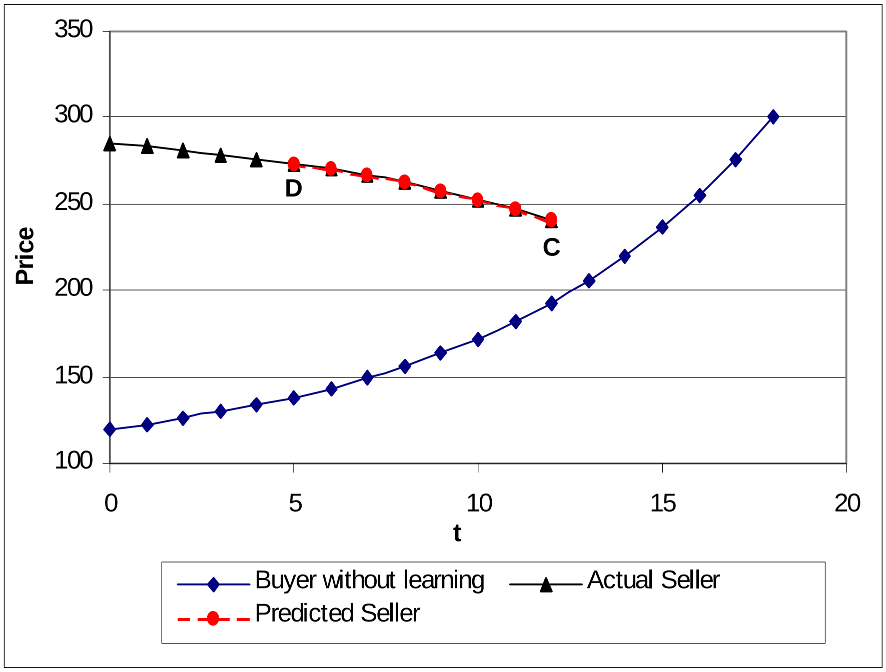

Figure 4.8. A Negotiation Process without any Possible Final Agreement.

4.5.2 Feasibility Check

After obtaining a range of estimated prices which are beneficial to an agent, we find out the best-estimated price that is attainable according to the parameters of an agent's own negotiation model. The best performance is achieved when the agreement is made at an opponent's reservation price, which represents the maximum (minimum) price at which a buyer (seller) allows an agent to conclude an agreement. Thus, we propose that each estimated price, starting from the reservation price, be justified for feasibility (Steps 2a and 2b). That is, each offer has to satisfy three criteria given below.

1) *The estimated price must be less than or equal to the agent's own reservation price.* Since a buying agent cannot accept a price that is greater than its reservation price, any estimated price greater than its reservation price is not feasible (Criterion 1 of Step 2c).
2) *The estimated price must be greater than the current bid price.* Theorem 1 states that the price of a buying agent increases with the number of turns. Thus, any estimated price less than or equal to the current bid price is not feasible. This can be observed in Figure 4.9 (See horizontal dotted line). The estimated prices at E and F are less than the current bid price at $t = 3$, when the agent learns at $t = 4$ (Criterion 2 of Step 2c).
3) *The turn number corresponding to the candidate price needs to be smaller than or equal to the agent's own* T_{max}. In Figure 4.9, the estimated prices at G and H satisfy the first two criteria, but do not satisfy this criterion (Criterion 3 of Step 2c). That is, the estimated price at J is the lowest offer that the buyer can hope to attain.

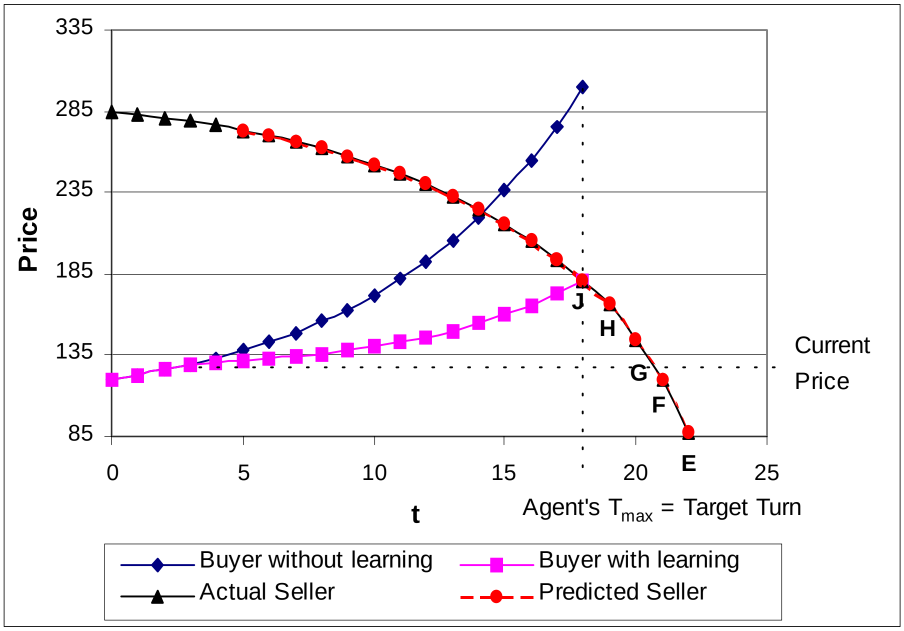

Figure 4.9. Effect of Learning on Negotiation Outcome (Target Turn = Agent's T_{max}).

4.5.3 Adjustment

The adjustment phase is executed if and only if an attainable price is found. In order to make the final price of the agreement to be the estimated one, an agent has to reach that desired price at the same turn as its opponent does. We can do this by adjusting the values of the parameters of the agent's time-dependent-tactical model. The parameters to be changed are dependent on the agent's own T_{max} and the target turn at which the estimated price is attained. Two cases result.

Case I: Target Turn = Agent's T_{max}

Figure 4.9 illustrates this case. As shown therein, the lowest attainable estimated price is at point J. With the target turn equal to the agent's T_{max}, point J is attainable by first adjusting the reservation price to equal J's price (Step 3a). Then, with the current β and K, the target turn and T_{max} are offset by the number of turns since the last adjustment (Step 3b). Finally, the price at the new $t = 0$ is reset to the current bid price, and the new P_{min} is computed by (4.1) as given in Step 3c. This sequence of changes ensures that J is reachable at the agent's T_{max}.

Case II: Target Turn < Agent's T_{max}

Figure 4.10 illustrates this case. The lowest attainable estimated price is at point L. Similar to *Case I*, we reset T_{max} and the target turn and compute the new P_{min} (Steps 3b and 3c). With the target turn smaller than the agent's T_{max}, we keep P_{max} unchanged, and compute the new β by (4.15), which is obtained by rearranging (3.1) (Step 3d).

$$\beta = \frac{\ln\left[\frac{\ln\left(\frac{\hat{P}(t_s)-P_{\min}}{P_{\max}-P_{\min}}\right)}{\ln(K)}\right]}{\ln\left(1-\frac{t_B}{T_{\max}}\right)} \tag{4.15}$$

where $\hat{P}(t_s)$ is the estimated price that an agent tries to reach, t_S, is the target turn, and t_B is the adjusted target turn. With the new parameters, the agent can reach L at the target turn.

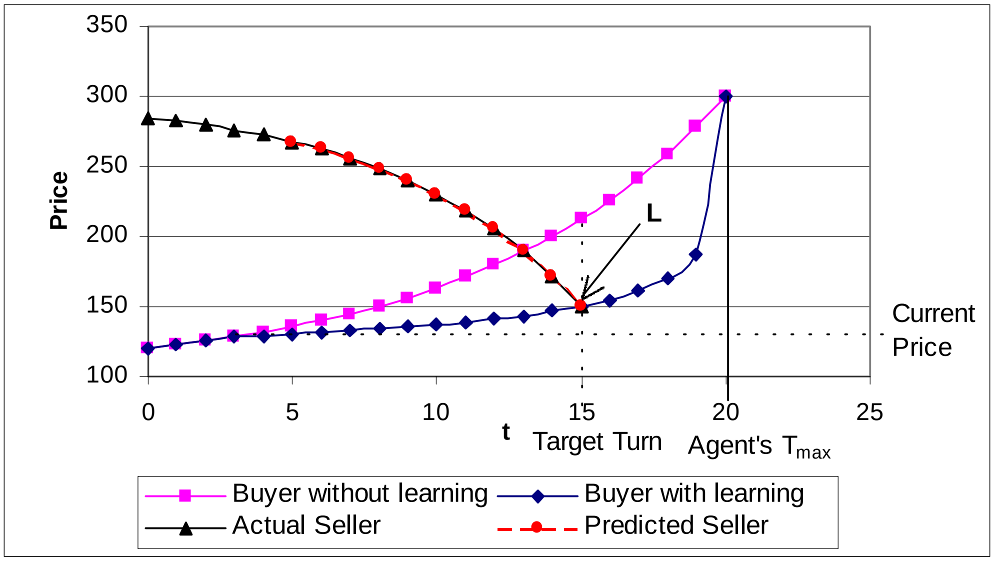

Figure 4.10. Effect of Learning on Negotiation Outcome (Target Turn < Agent's T_{max}).

1) Selection of Target Range

(a) Set $(\hat{P}_{\min}, \hat{P}_{\max}, \hat{T}_{\max}, \hat{\beta}, \hat{K})$ = the set of the estimated parameters from learning

(b) Set $t_{start} = \hat{T}_{\max}$

(c) Find out if there is a turn, t_i, for a possible agreement between the estimated model of the seller (estimated version of equation 3.2) and the agent's own model (equation 3.1)

(d) If t_i is present

(e) Set $t_{end} = t_i$

Else

(f) Set t_{end} = the turn at which the highest estimated price is located

End If

2) Feasibility Check

(a) For $t_S = t_{start}$ to t_{end} Step –1

(b) Find $\hat{P}(t_s)$ by (3.2); replace the parameters with $(\hat{P}_{\min}, \hat{P}_{\max}, \hat{T}_{\max}, \hat{\beta}, \hat{K})$, and t with t_S

(c) If $\hat{P}(t_s) \leq P_{\max}$ (Criteria 1)

and $P(t\text{-}1) < \hat{P}(t_s)$ (Criteria 2), where t = current turn

and $t_S \leq T_{\max}$ (Criteria 3), Exit For Loop and Goto Step 3

Next t_S

(d) Skip Step 3 when all $\hat{P}(t_s)$ are not attainable

3) Adjustment

(a) If $t_S = T_{\max}$

Set $P_{\max} = \hat{P}(t_s)$

End If

(b) Set $T_{\max} = T_{\max}$ – number of turns since last adjustment

Set $t_B = t_S$ – number of turns since last adjustment

(c) Find $P_{\min}$ by (4.1) and replace $P_{t=0}$ with $P(t\text{-}1)$

(d) If $t_B < T_{\max}$

Find β by (4.15)

End If

Note: This is the pseudo-code for buying agents. Modification is needed for selling agents.

Figure 4.11. Pseudo-code of the Reaction Algorithm.

4.6 Number of Offers

The learning algorithms discussed so far focus on learning using five price offers. However, agents are not guaranteed to receive five price offers from their opponents in the whole negotiation process. For example, when a selling agent starts a negotiation, the buying agent receives the first five price offers from the selling agent and estimates the parameters of the selling agent's negotiation model. Before the selling agent learns, the buying agent makes adjustment on the parameters of its negotiation model as given in §4.5. Clearly then, the selling agent cannot use five offers from the buying agent, because the last offer is generated by the new (adjusted) parameters. Hence, we propose that the selling agent estimates the parameters using four price offers from the buying agent. We call this as 4-point learning.

Although the buying agent obtains one more price offer for its first learning, this does not mean that the buying agent gains benefits in the whole negotiation process. Since the buying agent adjusts its parameters first at the end of the first round of learning, during the second round of learning, the seller receives five buyer's offers; by contrast the buyer has to do with four offers for the second round. Thus, an agent alternates between 4-point and 5-point learning, until an agreement is reached or the negotiation is terminated.

Both the learning and the reacting algorithms introduced in §4.4 and §4.5, respectively are applicable to 4-point and 5-point learning environments. The only difference is the value of N, which represents the number of price offers from an opponent. In 4-point learning environment, equation (3.1) at $t = 4$ is not available. Thus, the sum of squared error between the estimated and the actual prices consists of four error terms in 4-point learning environment, while that consists of five error terms in 5-point learning environment.

5 Experimental Study of Heuristics

In Chapter 4, we proposed heuristic algorithms for estimating the parameters of the time-dependent-tactical model. Our approach consists of three phases, and different algorithms were proposed in these phases. In this chapter, we report on simulation experiments to assess the effectiveness of these algorithms.

The performance measures used for testing are discussed in §5.1. After that, propositions are stated in §5.2, and the corresponding experimental designs are presented in §5.3. Experimental results and discussion appear in §5.4 and §5.5, respectively.

5.1 Performance Measure

Since the time-dependent-tactical model consists of five parameters, it is difficult to determine which estimation algorithm has better performance by comparing the estimated and the actual values of the individual parameters themselves. In order to simplify this, we use a pooled performance measure, the sum of squared error (*SSE*) between the actual price and the estimated price. This measure is also used in multiple regressions. One advantage of this performance measure is that it is computed by combining the values of all five actual and estimated parameters. We introduce three different types of *SSEs*. They differ from one another in the range of the turns that is used in the computation.

The first one restricts the *SSE* computation to the number of turns that has elapsed thus far in the negotiation process. It is denoted as SSE_F and is obtained by

$$SSE_F = \sum_{t=0}^{end}\left(P(t) - \hat{P}(t)\right)^2 \qquad \begin{cases} end = 3 & \text{for} \quad \text{4-point learning} \\ end = 4 & \text{for} \quad \text{5-point learning} \end{cases}$$

where $P(t)$ is the actual price, and $\hat{P}(t)$ is the estimated price as given in Chapter 4.

Another extreme is to compute the *SSE* over the entire negotiation time frame. However, since different estimation algorithms have different estimated T_{max} values that could in turn differ from the actual T_{max}, the number of error terms that must be included in the *SSE* computation becomes a concern. For example, when the actual T_{max} is greater than the estimated T_{max} by two turns, the estimated prices at t = actual T_{max} and t = actual $T_{max} - 1$ are missing. Thus, there is a difficulty in defining an equation to compute this *SSE*. Hence, this performance measure is not used.

The third measure, denoted as SSE_N, is computed over the next subsequent turns. As discussed in §4.6, agents continue learning process at periodic intervals. Hence, the accuracy of an estimation algorithm at the next few turns is comparatively more important than that of the other two measures, because the parameters of the opponent's negotiation model are estimated again at periodic intervals. SSE_N is computed over the next five turns, and is given by

$$SSE_N = \sum_{t=start+1}^{start+5} \left(P(t) - \hat{P}(t) \right)^2$$

where start is the number of turns elapsed thus far.

In our experiment, both SSE_F and SSE_N were measured. Since the SSE_N is comparatively more important than the SSE_F, only the SSE_N is presented and used as the basis to validate the hypotheses presented hereon. However, during the negotiation process, SSE_N is not available, and only SSE_F is available. Thus, the relationship between SSE_F and SSE_N is addressed in §5.6. The purpose is to justify if SSE_F could be a proxy measure for SSE_N.

5.2 Hypotheses

The hypotheses in this section are proposed based on our observation, when designing the heuristic learning algorithms.

5.2.1 Construction Phase

In the construction phase, the objective is to find a set of tentative estimates for reference, based on the derivatives at $t = 0$. Since accuracy is sacrificed while tractability is achieved, the accuracy of the estimated parameters in the construction phase is not expected to be good. According to §4.3, when the third-degree Taylor's polynomial is used, the accuracy of the estimated derivatives is lower for higher-order derivatives (See Table 4.1). By this, algorithms with higher-order derivatives can be expected to have lower accuracy. Thus, the two-derivative algorithm is predicted to have better performance than the three-derivative algorithm in the construction phase. We therefore propose the following hypotheses.

1) The *SSEs* of the construction-phase algorithms will be greater than zero.
2) The two-derivative algorithm will perform better than the three-derivative algorithm in the construction phase.

5.2.2 Improvement Phase

We propose three hypotheses for the improvement phase. First, the objective of the improvement phase is to obtain the best set of estimated parameters by defining a search range, based on the reference estimations obtained from the construction phase. Since we start with these estimates and enumerate further, predictability is expected to be enhanced in the improvement phase.

Our next hypothesis concerns the comparison among the three estimation algorithms in the improvement phase. These three algorithms differ from one another in the equations applied in the enumeration. Since the price offers in the price algorithm are actual values, while the derivatives in the other two algorithms are estimated from the best-fit curve, the price algorithm is predicted to outperform the other two algorithms.

Finally, as discussed in Chapter 3, the value of β determines the concession behavior in the time-dependent-tactical model. When the value of β is close to zero or too high, highly boulware or highly conceder behavior are expected. Based on our observation when designing the heuristic learning algorithms, fitting the price offers onto a best-fit curve for derivative estimation is more difficult when an opponent's behavior is highly boulware or highly conceder. Higher inaccuracy of the estimated derivatives is expected. Therefore, all three improvement-phase algorithms are expected to perform better for β's close to one than for other β's. Hence, the following hypotheses are suggested.

3) The predictability will be improved from the construction phase to the improvement phase.
4) The price algorithm will outperform the other two estimation algorithms in the improvement phase.
5) All three estimation algorithms in the improvement phase will perform better for $\beta \approx 1$ than for other β's.

5.3 Experimental Design

Experiment #1) This experiment is designed for testing hypotheses 1 to 4. There were two sets of replications in the experiment, one each for 5-point and 4-point learning. For each experimental set of replications, one thousand sets of parameters of the time-dependent-tactical model were randomly generated. The parameters that were held common to both sets were generated as given in Table 5.1. In this experiment, the objective is to test the general performance of our learning algorithms. In real-world cases, different users or agents usually have different values of the parameters in their negotiation models. To obtain an average performance measure, we vary the parameters as specified in Table 5.1.

Table 5.1. Common Parameters for Experiment #1.

Parameter	Range
β	U[0.1, 10]
T_{max}	U[20, 40]
K	U[0.1, 0.9]
P_{min}	U[100, 250]
P_{max}	U[300, 600]

where U[a, b] represents a uniform distribution with a range of a and b.

In both experimental sets, both construction-phase algorithms were run. As discussed in §4.4.2, depending on the accuracy of the estimated derivatives from the best-fit curve, a tentative set of parameters is not guaranteed to be found in the construction phase. Thus, when a solution was not available in both algorithms, learning was terminated and improvement phase was not started. If solutions were available in both algorithms, the estimated parameter set with the least SSE_F was selected as the final parameter set in the construction phase. If a solution was only available in one of the two estimation algorithms, the estimated parameters were selected from that algorithm. Given the construction-phase estimates, the improvement phase always yields a solution. Thus, all three improvement-phase algorithms were run.

Experiment #2) In order to test the interaction effect between the value of β and the estimation algorithm in the improvement phase (Hypothesis 5), a 3 x 3 factorial experimental design was established, with three β groups and the three improvement-phase algorithms as the two factors. The experiment consisted of three treatments, and each experimental treatment was carried out with the assigned β group according to Table 5.2.

Table 5.2. The Range of β in Each β Group in the Experiment #2.

β Group	**Lower Limit**	**Upper Limit**	**Experimental Treatment Number**
Low	0.01	0.5	1
Medium	0.75	1.25	2
High	5.0	7.5	3

$T_{max} \in$ U[20, 22], $K \in$ U[0.1, 0.2], $P_{min} \in$ U[100, 110], and $P_{max} \in$ U[250, 260]

In each experimental treatment, the values of the other four parameters were varied within a small range, as given at the bottom of Table 5.2. One thousand sets of parameters were randomly generated, and all three experimental treatments consisted of the same sets of parameters except for the β-values. The experiment was conducted under 5-point learning only.

5.4 Results

ANOVA tests and t-tests were conducted to test the hypotheses. The statistical test results for each of the five hypotheses are addressed in order.

5.4.1 Performance of Construction Phase

Hypothesis 1. (Estimates from Construction Phase are not Accurate) Table 5.3 presents the mean SSE_N from Experiment #1 for the construction-phase algorithms in both 5- and 4-point learning environments. Among 1000 sets of parameters in 5-point learning environment, 973 sets of estimated parameters were successfully found by the two-derivative algorithm, while 862 sets were found by the three-derivative algorithm. If the accuracy of the estimated parameters in the construction phase is good enough, the mean SSE_N will have insignificant differences from zero. Results from t-tests (Test value = 0, 95% confidence) indicate that the mean SSE_N of the two-derivative algorithm (M = 45.208) and the three-derivative algorithm (M = 747.959) are significantly different from zero ($p < 0.0005$).

In 4-point learning environment, 988 sets of estimated parameters were successfully found by the two-derivative algorithm, while 894 sets were found by the three-derivative algorithm. According to Table 5.3, results from t-tests (Test value = 0, 95% confidence) indicate that the mean SSE_N of the two-derivative algorithm (M = 39.311) and the three-derivative algorithm (M = 599.195) are significantly different from zero ($p < 0.0005$). In sum, hypothesis 1 is supported.

Table 5.3. Accuracy of the Construction-Phase Algorithms.

Algorithm	5-point Learning		4-point Learning	
	SSE_N	t-test p-value	SSE_N	t-test p-value
The Two-derivative Algorithm	45.208	< 0.0005	39.311	< 0.0005
The Three-derivative Algorithm	747.959	< 0.0005	599.195	< 0.0005

Note. t-tests were conducted with the test value = 0 at 95% confidence level

Hypothesis 2. (Two-derivative Algorithm outperforms Three-derivative Algorithm) This hypothesis is supported. A t-test[1] (95% confidence) between two construction-phase algorithms indicates that the mean SSE_N of the two-derivative algorithm is significantly lower than that of the three-derivative algorithm in both 5-point ($p < 0.0005$) and 4-point ($p < 0.0005$) learning environments.

5.4.2 Performance of Improvement Phase

Hypothesis 3. (Accuracy is gained in the Improvement-Phase Algorithms) Table 5.4 shows the mean SSE_N of the construction phase and three improvement-phase algorithms for both 5-point and 4-point learning in Experiment #1.

In 5-point learning environment, t-tests (95% confidence) between the construction phase and the improvement-phase algorithms indicate that the mean SSE_N of the price algorithm (M =

[1] Although there is a big difference between the two construction-phase algorithms, we formally conduct t-tests for the sake of uniformity and to make sure that all experimental results are statistically valid.

30.980), the first-derivative algorithm (*M* = 16.776) and the-second derivative algorithm (*M* = 19.418) in the improvement phase are significantly lower than that in the construction phase (*M* = 69.587, $p < 0.05$ for all comparisons). The hypothesis is supported in 5-point learning environment.

Table 5.4. Accuracy Gained by the Improvement-Phase Algorithms.

Learning Environment	Phase	Algorithm	SSE_N	t-test p-value
5-point Learning	Construction	Final Selected Parameter Set	69.587	--
	Improvement	The Price Algorithm	30.980	< 0.0005
		The First-Derivative Algorithm	16.776	0.004
		The Second-Derivative Algorithm	19.418	0.007
4-point Learning	Construction	Final Selected Parameter Set	42.054	--
	Improvement	The Price Algorithm	10.135	< 0.0005
		The First-Derivative Algorithm	17.293	< 0.0005
		The Second-Derivative Algorithm	34.466	0.265

Note. Final Selected Parameter Set indicates the final parameter set selected in the construction phase according to the minimum SSE_F between the two estimation algorithms.

In 4-point learning environment, similar results are obtained in both the price and the first-derivative algorithms. We conduct t-tests (95% confidence) between the construction phase and the improvement-phase algorithms, and find that the mean SSE_N of the price algorithm (*M* = 10.135) and the first-derivative algorithm (*M* = 17.293) in the improvement phase are significantly lower than that in the construction phase (*M* = 42.054, $p = 0.001$, $p < 0.0005$ respectively). However, the mean SSE_N of the second-derivative algorithm is not significantly different from that in the construction phase ($p = 0.265$).

Hypothesis 4. (Price Algorithm outperforms the other Two Improvement-Phase Algorithms) This hypothesis is partially supported. Table 5.5 shows the mean SSE_N and one-way ANOVA test results of three improvement-phase algorithms for both 5-point and 4-point learning in

Experiment #1. In addition, the post-hoc results of SSE_N of three algorithms for both 5-point and 4-point learning are illustrated in Table 5.6.

Table 5.5. Comparisons among the Improvement-Phase Algorithms.

Learning Environment	Algorithm	SSE_N	One-way ANOVA Test p-value
5-point Learning	The Price Algorithm The First-Derivative Algorithm The Second-Derivative Algorithm	30.980 16.776 19.418	0.379
4-point Learning	The Price Algorithm The First-Derivative Algorithm The Second-Derivative Algorithm	10.135 17.293 34.466	< 0.0005

Table 5.6. Post-hoc Results for comparing the Improvement-Phase Algorithms.

Learning Environment	Algorithms Comparison	Post-hoc p-value
5-point Learning	The Price Algorithm vs. The First-Derivative Algorithm The Price Algorithm vs. The Second-Derivative Algorithm The First-Derivative Algorithm vs. The Second-Derivative Algorithm	0.389 0.535 0.968
4-point Learning	The Price Algorithm vs. The First-Derivative Algorithm The Price Algorithm vs. The Second-Derivative Algorithm The First-Derivative Algorithm vs. The Second-Derivative Algorithm	0.197 < 0.0005 < 0.0005

In 5-point learning environment, when we conduct a one-way ANOVA test for the three improvement-phase algorithms, we do not observe any significant difference in the mean SSE_N among the algorithms. This is further illustrated by the post-hoc test results in Table 5.6. In 4-point learning environment, there is a significant difference. Post-hoc results show that the mean SSE_N of the price algorithm (M = 10.135) and the first-derivative algorithm (M = 17.293) are significantly lower than that of the second-derivative algorithm (M = 34.466, $p < 0.0005$, $p < 0.0005$ respectively), but we do not observe any significant difference in SSE_N between the price and the first-derivative algorithms.

In sum, the price algorithm and the first-derivative algorithm outperform the second-derivative algorithm in 4-point learning environment. On the other hand, in 5-point learning environment, no conclusion can be made to identify which one of the three improvement-phase algorithms has the best performance. This leads us to Experiment #2 for testing hypothesis 5.

Hypothesis 5. (Interaction Effect between β and Improvement-Phase Algorithms) A 3 x 3 two-way ANOVA test was conducted to investigate the interaction effect between β and the improvement-phase algorithm. Table 5.7 shows the mean SSE_N for different β and improvement-phase algorithm combination. Test result indicates that this hypothesis is partially supported.

Table 5.7. The Mean SSE_N for different β - Algorithm Combination.

Algorithm	β Group			
	Low*	Medium**	High***	All β Groups
The Price Algorithm	1.59×10^{-3}	1.80×10^{-2}	525.986	175.335
The First-Derivative Algorithm	5.19×10^{-3}	2.36×10^{-2}	37.565	12.531
The Second-Derivative Algorithm	1.23×10^{-3}	4.03×10^{-3}	2.221	0.742
All Algorithms	2.67×10^{-3}	1.52×10^{-2}	188.591	62.870

**Low $\beta \in [0.01, 0.5]$ **Medium $\beta \in [0.75, 1.25]$ ***High $\beta \in [5.0, 7.5]$*

The effect of improvement-phase algorithm on the SSE_N in different β groups is depicted in Figure 5.1, and the corresponding post-hoc test results are presented in Table 5.8.

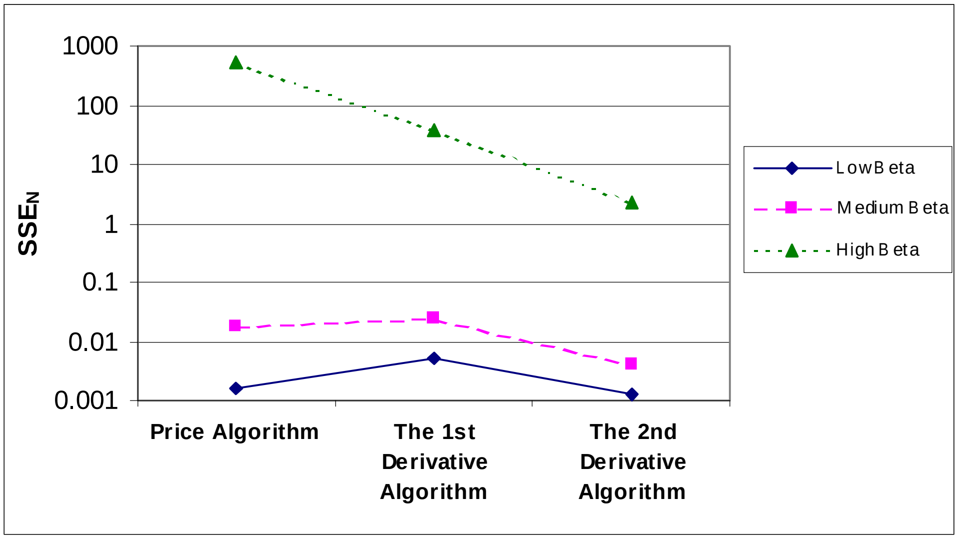

Figure 5.1. Effect of Improvement-Phase Algorithm on SSE_N in Different β Groups.

According to Figure 5.1 and Table 5.8, the mean SSE_N of the second-derivative algorithm (M = 1.23 x 10^{-3}) and the price algorithm (M = 1.59 x 10^{-3}) are significantly lower than that of the first-derivative algorithm (M = 5.19 x 10^{-3}), when only the result in low β group is considered. Significant difference between the price algorithm and the second-derivative algorithm is not found. In both medium and high β groups, we observe that the second-derivative algorithm outperforms the price algorithm and the first-derivative algorithm.

Table 5.8. Post-hoc Results for the Effect of β Group on the Three Improvement-Phase Algorithms.

β Group	Algorithm Comparison	Post-hoc p-value
Low	The Price Algorithm vs. The First-Derivative Algorithm The Price Algorithm vs. The Second-Derivative Algorithm The First-Derivative Algorithm vs. The Second-Derivative Algorithm	< 0.0005 0.531 < 0.0005
Medium	The Price Algorithm vs. The First-Derivative Algorithm The Price Algorithm vs. The Second-Derivative Algorithm The First-Derivative Algorithm vs. The Second-Derivative Algorithm	0.370 0.002 < 0.0005
High	The Price Algorithm vs. The First-Derivative Algorithm The Price Algorithm vs. The Second-Derivative Algorithm The First-Derivative Algorithm vs. The Second-Derivative Algorithm	< 0.0005 < 0.0005 < 0.0005

The effect of β group on the SSE_N in different improvement-phase algorithms is depicted in Figure 5.2, and the corresponding post-hoc test results are presented in Table 5.9.

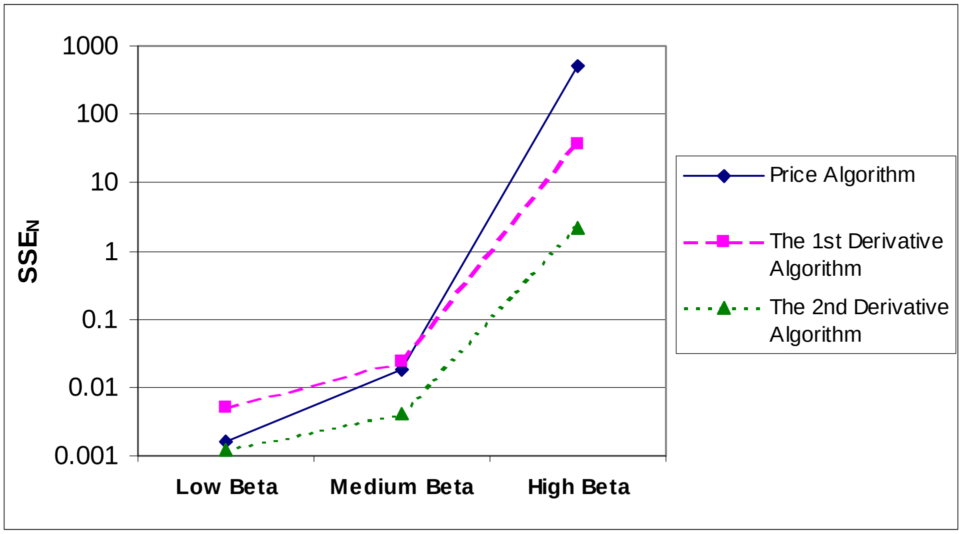

Figure 5.2. The Effect of β Group on SSE_N in Different Improvement-Phase Algorithms.

Table 5.9. Post-hoc Results for the Effect of Improvement-Phase Algorithm on β Groups.

Algorithm	**β Group Comparison**	**Post-hoc Test p-value**
The Price Algorithm	Low vs. Medium Low vs. High Medium vs. High	0.999 < 0.0005 < 0.0005
The First-Derivative Algorithm	Low vs. Medium Low vs. High Medium vs. High	0.999 < 0.0005 < 0.0005
The Second-Derivative Algorithm	Low vs. Medium Low vs. High Medium vs. High	0.999 < 0.0005 < 0.0005

According to Figure 5.2 and Table 5.9, for all three improvement-phase algorithms, the mean SSE_N of low and medium β groups are significantly lower than that of high β group. However, we do not find any significant difference between low and medium β groups. In summary, hypothesis 5 is partially supported.

5.5 Discussion

In sum, the estimation model given by the construction-phase algorithms is not accurate enough. This is expected because we focus on simplifying the equations rather than obtaining accurate estimates in the construction phase. In addition, the two-derivative algorithm is found to outperform the three-derivative algorithm in the construction phase. This is reasonable because the three-derivative algorithm consists of the estimated third derivative, which is not expected to be highly accurate. Thus, the two-derivative algorithm achieves better predictability.

The price algorithm, the first-derivative algorithm and the second-derivative algorithm are found to have improvement in predictability from the construction phase to the improvement phase in 5-point learning environment. However, only the price and the first-derivative algorithms have improvement in 4-point learning environment. According to Table 5.3, the second-derivative algorithm also demonstrates improvement in predictability from the construction phase to the improvement phase in 4-point learning environment, but the improvement is not significant enough. Generally, higher predictability is achieved in the improvement phase than in the construction phase. The result is expected because the improvement phase focuses on accuracy of the estimated parameters.

In the improvement phase, both the price algorithm and the first-derivative algorithm outperform the second derivative algorithm in 4-point learning environment. In 5-point learning environment, no conclusion can be made to identify which one of the three estimation algorithms has the best performance, based on the results in Experiment #1. However, according to the results in Experiment #2, the second-derivative algorithm is found to outperform the other two improvement-phase algorithms in both medium and high β groups in 5-point learning environment. In addition, the second-derivative and the pirce algorithms achieve higher performance than the first-derivative algorithm in low β group. Overall, the second-derivative algorithm is recommended to be implemented in 5-point learning environment, while either the price algorithm or the first-derivative algorithm is recommended in 4-point learning environment. Table 5.10 summarizes the algorithms used in the construction and the improvement phases.

Table 5.10. Summary of the Algorithms in the Construction and the Improvement Phases.

Phase	Learning Environment	Algorithm
Construction	4-point Learning 5-point Learning	The Two-Derivative Algorithm
Improvement	4-point Learning 5-point Learning	The Price Algorithm or The First-Derivative Algorithm The Second-Derivative Algorithm

Finally, all three estimation algorithms in the improvement phase are found to have lower predictability in cases where β is high than in cases where β is low or close to one. That is, based on our tests, our learning algorithms are more effective on agents with boulware or linear behaviors than with conceder behavior.

5.6 Proxy Measure for SSE_N

Our results so far are based on the performance measure of SSE_N, because, as discussed in §5.2, the SSE_N is comparatively more important than SSE_F. However, the computation of SSE_N involves un-received price offers. Hence, during the negotiation process, we can get SSE_F but not SSE_N. The question then is whether SSE_F could be a proxy measure for SSE_N.

To test this, we determine the correlation coefficient between the SSE_F and the SSE_N. Prediction of the SSE_N based on the SSE_F was tried unsuccessfully. Table 5.11 gives the correlation coefficient between the SSE_F and the SSE_N for our various algorithms in Experiment #1.

Table 5.11. Correlation Coefficient between the SSE_F and the SSE_N for various Algorithms.

Learning Environment	Phase	Algorithm	Correlation Coefficient	p-value
5-point Learning	Construction	The Two-derivative Algorithm	0.451	< 0.0005
		The Three-derivative Algorithm	0.508	< 0.0005
	Improvement	The Price Algorithm	0.976	< 0.0005
		The First-Derivative Algorithm	0.315	< 0.0005
		The Second-Derivative Algorithm	0.127	< 0.0005
4-point Learning	Construction	The Two-derivative Algorithm	0.342	< 0.0005
		The Three-derivative Algorithm	0.540	< 0.0005
	Improvement	The Price Algorithm	0.774	< 0.0005
		The First-Derivative Algorithm	0.151	< 0.0005
		The Second-Derivative Algorithm	0.242	< 0.0005

Note. The algorithms in the shaded cells are the selected algorithms in the construction and the improvement phases of the proposed learning approach.

The results therein indicate that all correlation coefficients are significant. The SSE_F and the SSE_N are highly correlated for the price algorithm, but this does not appear to be so in other algorithms. Even though the correlation coefficients are low in some cases, we still use SSE_F as a proxy, because of two reasons. First, SSE_N is not available during a negotiation process. In addition, according to our experimental results in Chapters 8 and 9, when we use SSE_F as a proxy, our learning algorithms are effective in arriving at a solution, and more importantly, in improving our negotiation outcomes (e.g. individual utility, joint utility etc.).

6 Electronic Negotiation Agent Design

In Chapter 4, we have developed heuristic learning algorithms for estimating the parameters of the time-dependent-tactical model. In Chapter 5, we have conducted simulation to assess the effectiveness of the estimation algorithms in each phase. In this chapter, we introduce how to integrate our proposed learning algorithms into negotiation agents.

As discussed in Chapter 2, many electronic negotiation agents have been designed to test the feasibility of different negotiation models (Faratin *et al.*, 1997; Zeng & Sycara, 1998; Sim & Chan, 2000; Zacharia *et al.*, 2001; Deveaux *et al.*, 2001). A basic design of a negotiation agent consists of two steps: user interaction and negotiation process.

In the user interaction step, users specify their preferences to the agents. Common parameters include the minimum and maximum values of the negotiating issues. Further, depending on the agent design, other parameters include the time limit of the negotiation, the weight of each negotiating issue in the utility calculation, and the choice of negotiation tactics. In any case, once users' preferences are specified, agents start negotiations without further user interactions until negotiations end. The second part of an agent design is the negotiating process, in which buying and selling agents exchange offers. Agents evaluate the opponents' offers and make decisions to accept or reject the offers. If agents prefer to continue the negotiation process, they employ their negotiation tactics to determine the values of the negotiating issues, and make counter-offers. Agents continue the negotiation process until the negotiation ends, which occurs when either agent terminates the process because of time expiry, or when both parties compromise the values of the negotiating issues.

Sim & Chan (2000) have further modified the above design by adding a matching process that connects certain buying agents with certain other selling agents. At the beginning of the matching process, an agent looks for candidate opponents with whom to negotiate in the market (*Selecting*). After opponents are selected, the utility of each tentative connection is evaluated, and compared with a defined criterion (*Evaluating Connection*). If the utility exceeds certain threshold level, which will be referred to as the *filtering threshold* in this thesis, the connection between the agents is selected as one feasible negotiation. If the utility does not exceed this *filtering threshold*, the negotiation will not start (*Filtering Connection*). Agents continue this matching process until all possible connections are processed, or the number of feasible negotiations reaches certain limit. The flowchart of the electronic negotiation agent design containing the users' interaction with agents, the matching process, and the negotiation process is depicted in Figure 6.1.

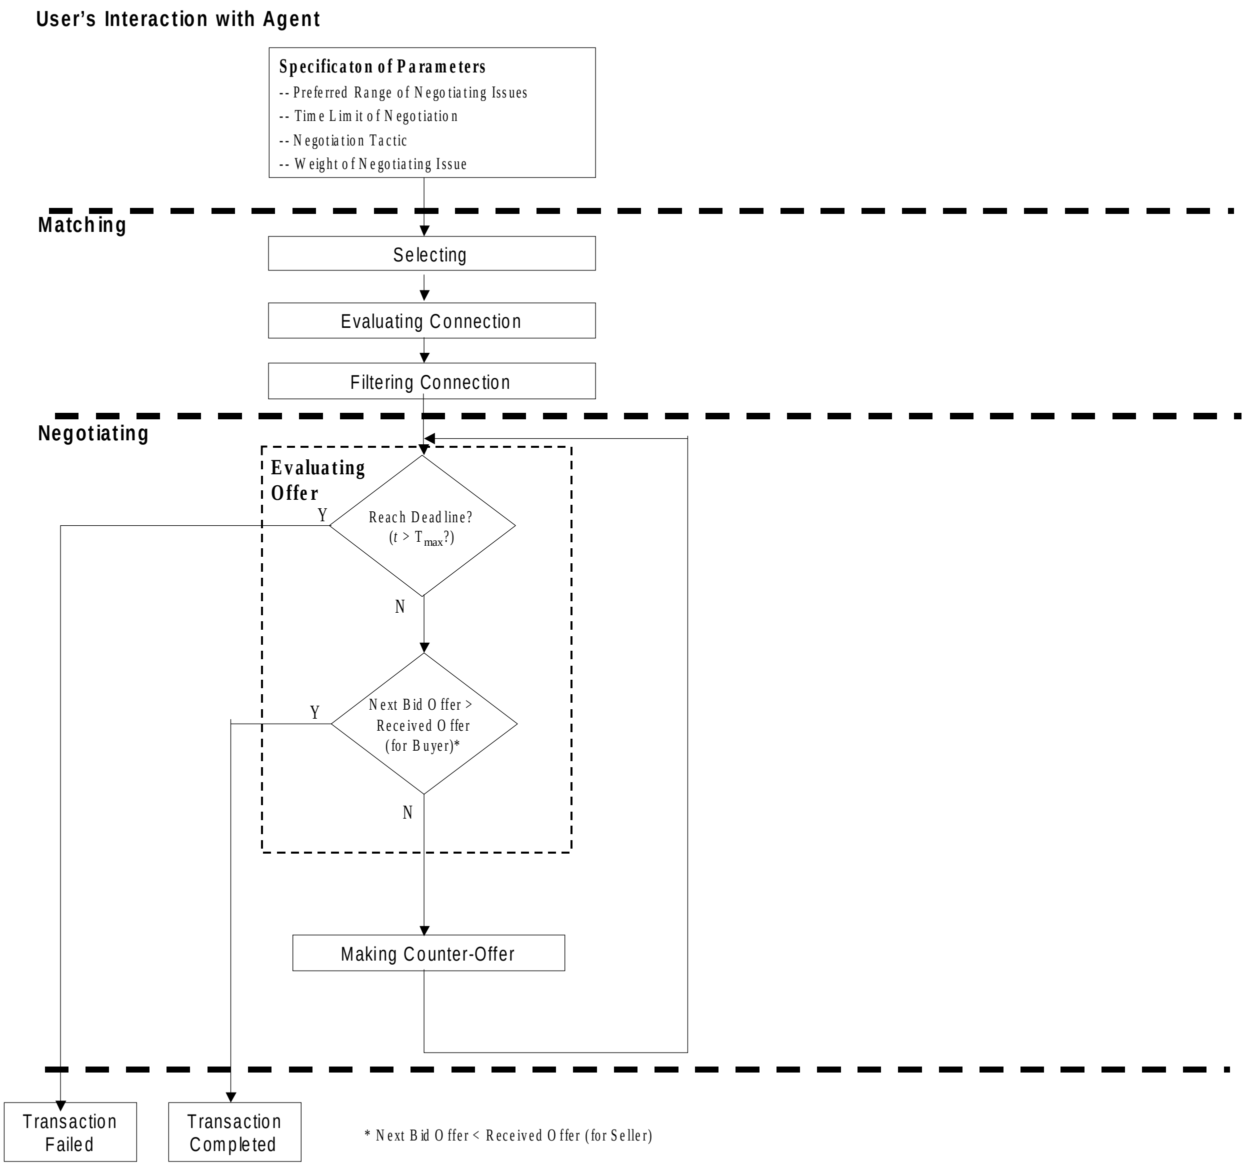

Figure 6.1. The Flowchart of the Electronic Negotiation Agent Design.

With Figure 6.1 as a framework, Sundarraj (2002) has introduced several new components into the design of a negotiation agent. These components include analytical hierarchy process (AHP), situational power, and goal constraints. In §6.1, we introduce how to integrate the analytical hierarchy process for computing the utility. In order to test the effects of situational power and goal constraints on electronic negotiation outcome, we present how to incorporate the situational power component as well as the goal setting component into our electronic negotiation agent. Finally, in §6.2, we show how the learning algorithms introduced in Chapter 4 can be integrated into our agent.

6.1 Integration of AHP, Situational Power and Goal Setting

The parameters to be specified into our electronic negotiation agent are identified in Figure 6.2. They are similar to those in Figure 6.1 except for some new parameters which are shaded. The negotiation tactic is the time-dependent-tactical model introduced in Chapter 3. The parameters of the time-dependent-tactical model represent the negotiating range of the negotiating issues, the time limit of the negotiation, and the type of negotiation behavior. Instead of specifying the weights of the negotiating issues for computing the utility, we use AHP to represent the importance level between any two negotiating issues according to the users' preference level. Situational power, which represents market conditions, is manipulated by AHP as well. Finally, we manipulate goal constraints by a utility level that users want their agents to accomplish in negotiations.

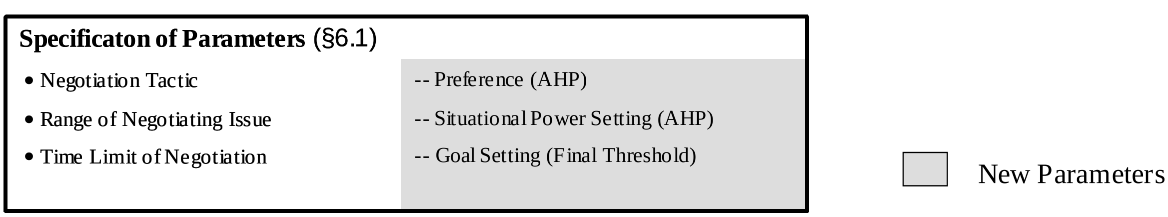

Figure 6.2. The Parameters to be Specified to the Negotiation Agent.

6.1.1 Analytical Hierarchy Process (AHP)

We reviewed different utility-based functions in Chapter 2, but AHP has not been implemented in any electronic negotiation agent so far. A negotiation often involves multiple attributes. In order to find an overall measure of the values of the negotiating attributes, utility can be evaluated by AHP, if users' preference levels of these attributes are specified into their negotiation agents (Sundarraj, 2002). In AHP, users' preference levels are specified in preference pairwise matrix to reflect their interests in different factors and properties. Besides the negotiating attributes, Sundarraj (2002) has also applied AHP to non-negotiating attributes. Non-negotiating attributes refer to those issues or factors that cannot be negotiated, but act as major components of the utility. For example, when a buyer wants to buy a house, and several houses are available from several different sellers, the locations of these houses are not negotiable, but they are major decision factors when the buyer makes choices.

There is a difference in the content of the preference pairwise matrix between a buying agent and a selling agent. The preference pairwise matrix of a buying agent involves both negotiating and non-negotiating attributes, while that of a selling agent consists of negotiating issues only. The reason is that non-negotiating issues are usually the properties of the selling products. Since the values of these issues are fixed, and are not considered as decision factors for selling agents in negotiations, they are not taken into account in the utility calculation for selling agents. Thus, for selling agents, the preference levels of the non-negotiating issues are not required in our design.

6.1.2 Modeling Situational Power

The situational power of an agent is defined as its perception about the preferences of the market at large for the product of interest. In our agent design, situational power represents the perceived power of an agent. Consistent with the method to elicit a buyer's preference in §6.1.1, situational power is also manipulated by AHP. This manner of manipulating situational power using AHP is similar to the application of AHP in benchmarking an organization's performance (Forman & Gass, 2001). By using AHP, an organization can compare itself with the best-of-breed organization, and thereby evaluating its process. In other words, a comparison of the AHP scores between two companies can indicate how much a company performs better than the other one does. We apply this idea to the context of situational power by comparing an agent's perception of the market preferences for the product of interest with a neutral set of preferences. That is, we compare the AHP utilities between buyers in the market with a neutral buyer, and use the difference to represent the level of power.

To perform the above comparison, we use preference pairwise matrix to specify the preferences of the market. Given an agent's perception of the market preferences, a *market utility*, *MU*, can be computed using AHP for each seller's product. To model situational power, we establish a neutral buyer who prefers all negotiating and non-negotiating issues equally. In other words, each cell in the preference pairwise matrix is equal to one. According to the preference of this neutral buyer, a *neutral utility*, *NU*, can also be computed by AHP. The difference between *MU* and *NU* is used as a measure of situational power. For example, when *MU* > *NU*, buyers in the market have comparatively a higher preference for the seller's product. In other words, there is a high competition from other buying agents, and the situational power of the buying agent is reduced. On the other hand, the power of the selling agent increases, because the market likes its product.

Pruitt (1981) and Rubin & Brown (1975) suggested that negotiators who see themselves as powerful usually extract a large concession from their opponents, and make small concession of their own. Mathematically, we react to this situational power by adjusting the value of β in the time-dependent-tactical model, as given below.

Buying agent: $$\text{New } \beta = \text{Old } \beta \times \left(1 + \min\left\{\frac{MU - NU}{NU}, \text{User - Specified } \beta \text{ Change}\right\}\right)$$

Selling agent: $$\text{New } \beta = \text{Old } \beta \times \left(1 - \min\left\{\frac{MU - NU}{NU}, \text{User - Specified } \beta \text{ Change}\right\}\right)$$

where User-Specified β Change represents the maximum allowable change of β specified by users. As *MU* is greater than *NU*, buyers in the market have a higher preference on the seller's product. Thus, buying agents see themselves as less powerful. In order to have higher competitive advantage, buying agents increase their concession rates. On the other hand, selling agents are said to have comparatively higher power, and they decrease their concession rates. When *MU* is smaller than *NU*, buying agents decrease their concession rates, while selling agents do the opposite.

6.1.3 Modeling Goal Setting

A goal for an agent can be represented by a utility value that a user wants the agent to achieve before concluding a negotiation. This utility will be named as *final threshold* in this thesis. When a *final threshold* is specified, an agreement can only be made when the individual utility of the agreement is greater than or equal to the *final threshold*.

Since users are not expected to perform utility calculation before specifying their *final threshold* values to agents, it is possible that the values of the negotiating issues in some of the offers do not satisfy the *final thresholds*. To prevent this anomaly, we adjust the reservation values of the negotiating issues. For example, when a buying agent negotiates the price of a product, the individual utility decreases from P_{min} to P_{max}. When a *final threshold* is specified, the utility at P_{max} might be lower than the *final threshold*. If this is the case, we use AHP to compute the price offer corresponding to the *final threshold*, and use that price as the new reservation price.

6.2 Integration of Learning

The algorithms to estimate the parameters of an opponent's time-dependent-tactical model, and the algorithm to react to the estimates have been presented in §4.4 and §4.5, respectively. In this section, we introduce how to integrate these algorithms into the agent outlined in §6.1.

When an agent reaches the turn to learn as in 4-point and 5-point learning discussed in §4.6, it carries out the learning process as described below. The learning component in our negotiation agent consists of three phases: (1) Checking, (2) Estimating and (3) Reacting. The flowchart of these three phases is depicted in Figure 6.3. The objective of the checking phase is to validate if an opponent's offers since the last learning still satisfy the current estimation model that has been constructed at the last learning turn. If an opponent's moves since last learning cannot be predicted by the current estimation model, an agent obtains a new set of estimated parameters in the estimating phase, and adjusts the parameters of its negotiation model in the reacting phase. On the other hand, if the current estimation model is still valid, the estimating and the reacting phases are skipped. The validity of the current estimation model is justified, based on the value of the regression parameter, R^2. If $R^2 \geq 0.95$, we assume that the agent continues using the current estimation model for prediction.

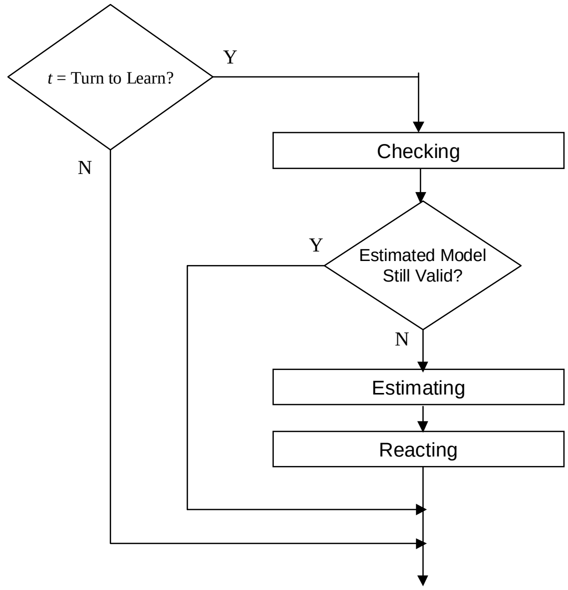

Figure 6.3. The Flowchart of Learning Component of Negotiation Agent Design.

In the estimating phase, an agent implements the algorithms in §4.4 to determine the parameters of its opponent's time-dependent-tactical model. According to the results and discussion in Chapter 5, our negotiation agent carries out the two-derivative algorithm in the construction phase. In the improvement phase, we implement the second-derivative algorithm and the price algorithm in 5-point and 4-point learning environments, respectively.

In the reacting phase, an agent implements the reaction algorithm in §4.5 to adjust the parameters of its time-dependent-tactical model.

6.3 Summary

When the power setting, goal setting and learning are integrated into the agent depicted in Figure 6.1, a revised flowchart of the negotiation agent design can be presented as in Figure 6.4. First, a user specifies his/her preferences to an agent. In the matching step, the agent looks for opponents for tentative negotiations in the market. When candidate opponents are found, the agent adjusts the parameters of its negotiation tactic according to the perceived power and goal setting. Then the agent computes the utilities for all potential connections, and decides which connections are feasible based on the *filtering threshold*. After that, the agent proceeds to the negotiating step, and exchanges offer with its opponents. There are three stopping criteria to the negotiation process. First, if the agent reaches its time limit, the negotiation is terminated, and becomes failed. Second, if the agent has a goal constraint, and reaches its last negotiation turn, it completes the transaction when the received offer satisfies the goal. When these two criteria are not satisfied, the agent computes its next offer, and compares with the current received offer. For a buying (selling) agent, if the utility function of the negotiating issue is decreasing (increasing), and its next offer is greater (less) than the current received one, an agreement is made based on the received offer. Otherwise, the agent keeps on negotiating by making a counter offer. This negotiating process continues until either party terminates the process, or both agents compromise for an agreement. When the agent learns, it implements the learning algorithms at periodic intervals. The agent carries out learning before it computes its next offer. This ensures that the next offer is built based on the new negotiation model. The agent stops working when all feasible negotiations are concluded.

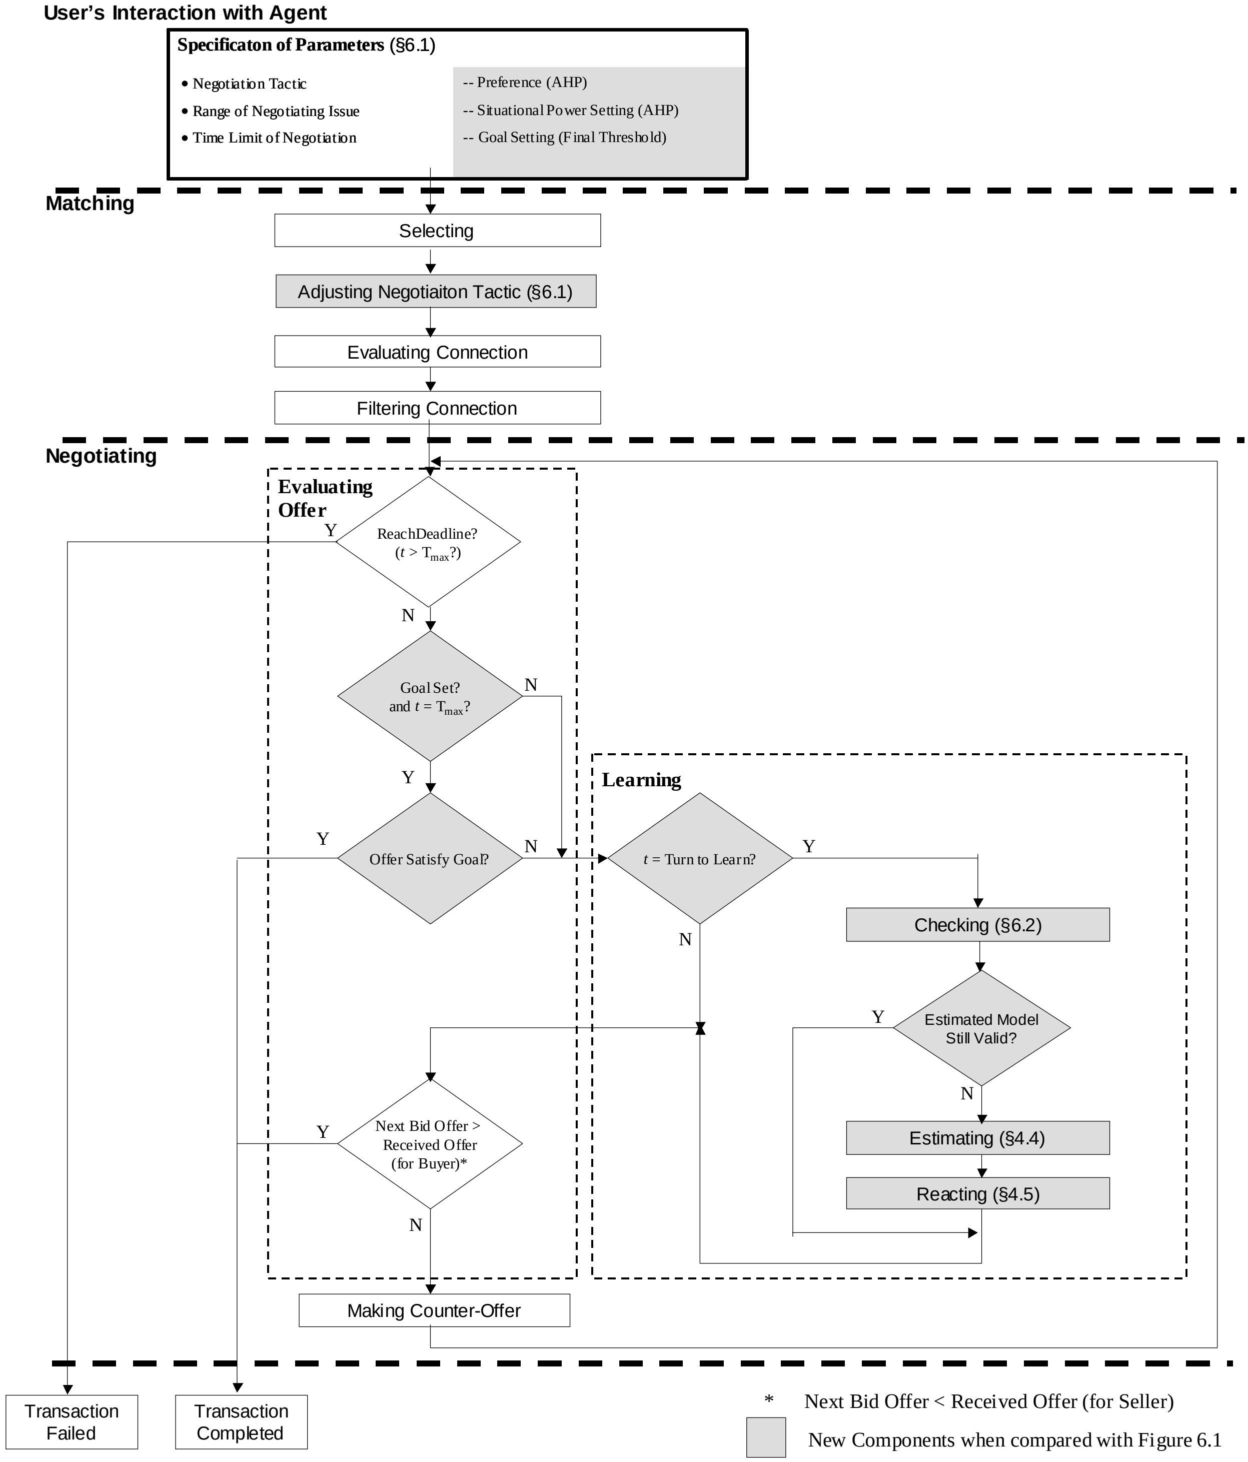

Figure 6.4. The Revised Flowchart of the Electronic Negotiation Agent Design.

7 Hypotheses from Human Negotiation Research

Thus far in this thesis, we have presented how situational power, goal setting, and learning can be incorporated into our electronic negotiation framework. In this chapter, we present the hypotheses on the effects of situational power, goal setting and learning on human negotiation outcome. Also, we present the experimental design for testing the hypotheses in the electronic-negotiation context. Results with these experiments will be given in the next two chapters.

The effects of perceived power and goal setting on negotiation outcome have been frequently studied in the past few decades. Several hypotheses have been tested and verified in human negotiation. However, the effects of power, goal setting, and learning on electronic negotiation have not been established. Since one desirable goal in this thesis is to make electronic negotiations to be more human-like, the agreement becomes an important element to assess the effectiveness of our integration methodologies. In order to fill in this gap, hypotheses tested in human negotiation are taken up for validation in the electronic setting. Our objective here is not to propose new theories, but to test the level of agreement between human and electronic negotiations. Although the effect of learning has been tested by incorporating Bayesian learning and genetic algorithms into electronic negotiation agents, the joint effect of learning and power as well as that of learning and goal setting have not been examined before. Several hypotheses on joint effects are proposed and tested in this thesis.

First, performance measures used for validation are identified in §7.1. Then the effects of perceived power and goal setting in human negotiation are reviewed. Also, hypotheses on the effects of perceived power, goal setting as well as the joint effects by incorporating learning in electronic negotiation are proposed. The experimental design is given in §7.5.

7.1 Performance Measure

We introduce six performance measures to validate the hypotheses proposed in this chapter. A summary of the measures is presented in Table 7.1.

Table 7.1. Summary of Performance Measures.

Name of Performance Measure	Notation of Performance Measure
Impasse Rate	*IR*
Individual Utility	*IU*
Joint Utility	*JU*
Number of Turns	*NT*
Distance to Nash Solution	*DN*

Impasse Rate (IR)

The impasse rate of an agent represents the failure rate of negotiations that have been started by the agent. Impasse rate is the number of incomplete transactions divided by the number of transactions in which the agent has participated.

Individual Utility (IU)

The individual buying (selling) agent utility represents the utility that a buying (selling) agent achieves in a dyad, when an agreement is made. It is computed by the AHP technique outlined in Chapter 6, and is based on the final values of the negotiating issues in a negotiation as well as on the preference pairwise matrices of the agent. According to Tripp & Sondak (1992), we take into account the impasse rate in utility computation, or else, experimental results might be biased. When a dyad fails to reach an agreement, an impasse is declared, and both the buying and selling agents receive zero utilities for that negotiation (Hamner and Harnett, 1974; Pinkley *et al.*, 1994).

Joint Utility (JU)

The joint utility of a dyad (i, j), which consists of a buying agent i and a selling agent j, is the sum of the individual utility of the buying agent and the individual utility of the selling agent.

$$JU_{ij} = BU_{ij} + SU_{ij} \qquad (7.1)$$

where BU is the buying agent's IU, and SU is the selling agent's IU.

Number of Turns (NT)

This measure indicates the number of turns needed by a dyad to successfully complete an agreement, or the number of turns that has elapsed before the negotiation is terminated.

Distance to Nash Solution (DN)

The distance to Nash solution measures the fairness of a completed negotiation (Goh *et al.*, 2000; Lim *et al.*, 1993). The Nash solution is the settlement in which the product of the utilities of the buyer and the seller is maximized. According to Goh *et al.* (2000) and Lim *et al.* (1993), the distance to Nash solution of a dyad (*i*, *j*) is computed as

$$DN = \sqrt{\left(BU_{ij} - BN_{ij}\right)^2 + \left(SU_{ij} - SN_{ij}\right)^2} \tag{7.2}$$

where *BN* and *SN* represent the buying agent's and the selling agent's utilities corresponding to the Nash solution.

In this thesis, we will use the above performance measures in the manner that they were used in the human negotiations described below. Thus, only a subset of the performance measures is used in each experiment.

7.2 Effect of Perceived Power

As discussed in Chapter 6, situational power is a function of market-level behavior. Understanding market condition is one important factor to obtain competitive advantage. Zacharia *et al.* (2001) investigated the microeconomic effects of a dynamic pricing algorithm in a reputation-brokered agent. They observed that this dynamic pricing algorithm was found to be effective based on market conditions. Part *et al.* (1999) also stated that modeling the overall market behavior of an auction process is more effective than modeling the interior reasoning of each agent participating in an auction, because modeling the negotiation behavior of each agent is too enormous and takes time. Thus, the incorporation of market conditions into electronic negotiation agents becomes important, although the above-mentioned researchers did not incorporate power.

In the field of psychology, power is defined as the level of dependency between two parties. Power of A over B can be represented by the dependency of B on A for certain things (Emerson, 1962). In the past few decades, the effect of power on human negotiation has been frequently studied and tested. Pinkley *et al.* (1994) designed an experiment which consisted of buyers and sellers with high level of BATNA (Best Alternative to The Negotiated Agreement), low level of BATNA and no BATNA. A 3 x 3 experiment was conducted by allowing all possible combinations of subjects to negotiate with each other in a fixed time period. McAlister *et al.* (1986) investigated the effect of power by setting up an experiment with four market conditions. The market condition is dependent on the ratio of the number of seller to the number of buyer, and on the maximum number of transactions that buyers and sellers were allowed to complete. Dwyer & Walker (1981) conducted an experiment, which consisted of balanced power in one condition and imbalanced power in another one, again by varying the number of negotiators in the market. The experiment consisted of one manufacturer and one retailer in the balanced power condition, and one manufacturer and two retailers in the imbalanced power condition. Thus, the manufacturer gains situational power in the latter condition. Greenhalgh *et al.* (1985) investigated the joint effects of preferences, personality, and situational power on negotiation outcomes. In their experiment, situational power was manipulated by varying the number of alternative parties with which the subjects would meet their needs. Among these papers, several similar hypotheses were validated, although there are some differences as discussed below.

Individual Performance

The hypothesis that has been frequently tested is that negotiators with higher power outperform those with lower power. This hypothesis is strongly supported in all the studies. Pinkley *et al.* (1994) found that negotiators with high BATNAs would outperform those with low and no BATNAs, because they could earn higher utility in the negotiated agreements. McAlister *et al.* (1986) found that higher-power negotiators achieved agreements with higher profitability per transaction as well as greater overall profitability in the entire market than lower-power negotiators. Similarly, Dwyer & Walker (1981) found that the manufacturer gained more profits when the manufacturer gained power in imbalanced power situation. The reason for this finding is that negotiators with higher power tend to behave exploitatively, while those with lower power tend to behave submissively. As a result, the negotiated agreements tend to favor the high-power negotiators.

Joint Performance

Another hypothesis that receives high attention is that agreements will be more integrative in equal-power situations than in unequal-power situations. Both Dwyer & Walker (1981), and McAlister *et al.* (1986) obtained results to support that joint utility in an equal-power condition is higher than that in an unequal-power condition. Since the high-power negotiators tend to push for agreements that favor them on both high-priority and low-priority issues, the low-power negotiators lose advantage of trade-offs between differently prioritized issues. Thus, agreements become less integrative. On the other hand, Pinkley *et al.* (1994) proposed that joint utility in an imbalanced power situation is higher than that in a balanced power situation. In an imbalanced power situation, lower-power negotiators tend to increase their profits by finding a settlement that increases the size of the share of the resource pool of the higher-power negotiators, while simultaneously the lower-power negotiators keep part of the expanded pool. However, the experimental results from Pinkley *et al.* (1994) did not indicate any significant difference in joint utilities between balanced power and imbalanced power situations. Among 28 studies that addressed these two propositions, four supported Pinkley's proposition (See Rubin & Brown, 1975 for a summary), while the remaining 24 studies supported the proposition from Dwyer & Walker (1981) and McAlister *et al.* (1986).

Number of Turns

Dwyer & Walker (1981) investigated the effect of power on the number of offers exchanged, before successfully concluding a negotiation. The results showed that fewer number of offers were required in the imbalanced power situation than in the balanced power situation. In their experiment, the balanced power condition consisted of one manufacturer and one retailer, while the imbalanced power condition consisted of one manufacturer and two retailers. Thus, the retailers in the imbalanced power condition were considered to have lower power. Since the low-power retailers had to compete with each other to make successful deals with the sole high-power manufacturer, the low-power retailers were more anxious than the high-power manufacturer to complete the transaction. The number of bid offers to complete the negotiation is expected to be reduced.

Nash Solution

The effect of power on the Nash solution (Nash, 1950; Nash, 1953) has not been verified in any papers addressed above. Neslin & Greenhalgh (1983) and Eliashberg *et al*. (1986) found that Nash solution can be used to predict the outcome of a negotiation. This brings out the importance of the Nash solution, and hence, in this thesis, we employ the distance to Nash solution to measure the fairness of the negotiation agreements. Negotiation agreements with balanced power are expected to be fairer than those with imbalanced power. This hypothesis will be validated in electronic negotiation.

A summary of the hypothesis is as follows.

Individual Level

P1) Agents with higher perceived power will have higher individual utilities than those with lower perceived power.

Dyadic Level

P2) Joint utility in a balanced power condition will be higher than that in an imbalanced power condition.

P3) The levels of perceived power of both agents will have a positive and additive effect on the time to reach agreements.

P4) Distance to Nash solution in a balanced power condition will be lower than that in an imbalanced power condition.

7.3 Joint Effect of Perceived Power and Learning

The joint effect of power and learning has not been studied so far in both human and electronic negotiations. Zeng & Sycara (1998) found that learning is beneficial in a negotiation model. Deveaux *et al*. (2001) also found that an agent's success in a negotiation can be enhanced by adapting the agent's tactics to the behavior of the other party, based on its belief of its opponent. Thus, when agents learn, they are predicted to have better performance. Higher individual utility and higher joint utility are expected. In addition, power and learning are expected to have an interaction effect. Agents with higher power tend to make the agreements more favorable to them. If they learn, they might further pull the agreements to favor them with a greater extent.

On the other hand, low-power agents with learning tend to gain higher individual utilities than those without learning, although they simultaneously behave submissively based on the market condition. Learning is believed to be more effective on high-power agents than on low-power agents.

Since high-power agents are expected to achieve higher individual utilities than low-power agents, and learning is expected to be beneficial to the agents, we propose that when agents learn, agents with higher power still outperform those with lower power.

In addition, Zeng & Sycara (1998) investigated the effect of Bayesian learning on Nash solution. A Nash solution is the highest when both parties learn while is the lowest when only one party learns. A negotiation agreement becomes fairer as both parties learn.

Based on the discussion, the following hypotheses are proposed and will be tested on both the individual and dyadic levels.

Individual Level

PL1) Agents with learning will have higher individual utilities than those without learning.

PL2) Learning will provide a higher increase in individual utility for agents with higher power as compared to agents with lower power.

PL3) When learning is implemented, agents with higher power will have higher individual utilities than those with lower power.

Dyadic Level

PL4) In both balanced and imbalanced power conditions, agreements with learning will have higher joint utilities, and lower distances to Nash solution than those without learning.

7.4 Effect of Goal Setting

Besides power, high attention has also been focused on the effect of goal setting on human negotiation. Locke *et al.* (1981) stated that a goal is what an individual tries to accomplish. Based on a review of a number of laboratory and experimental studies on the effects of goal setting on the performance of a task, Locke *et al.* (1981) concluded that in 90% of the studies, higher performance is obtained with specific and challenging goals rather than with easy goals, "do your best" goals, or no goals. In order to test the effect of goal setting in human negotiation, further studies and experiments on integrative negotiation were conducted.

Individual Utility

According to Locke, negotiators with specific and difficult goals are expected to achieve higher individual utilities than those with easy or no goals. Hamner & Harnett (1974) defined two levels of goal setting as high and low for negotiators to negotiate within a fixed time period. They found that the higher a negotiator's goal, the better would be his/her performance on the bargaining task. Other studies showed the goal manipulation by setting minimum profit levels that negotiators had to obtain based on the payoff tables (Bazerman & Neale, 1985; Huber & Neale, 1986; Huber & Neale, 1987). Participants were asked not to make agreements when the minimum profit level cannot be reached. Bazerman & Neale (1985) found that negotiators with moderately difficult set profit constraints would achieve more profitable individual agreements and greater overall profitability than negotiators without set constraints. Both Huber & Neale (1986) and Huber & Neale (1987) found a positive and linear relationship between goal difficulty and a negotiator's performance. In addition, Huber & Neale (1986) indicated that individual negotiator's performance would be higher for specific goal than for non-specific goal. These findings will be verified in electronic negotiation in this thesis.

Pruitt (1981) stated that negotiators with a higher level of goal difficulty might be more likely to despair of reaching agreements, because they are further apart from each other at any given time. However, this has not been validated in any of the papers cited above. The effect of goal setting on the impasse rate of a negotiation will be tested. Hypotheses on the relationship between goal difficulty and impasse rate as well as that between goal difficulty and individual negotiator's performance are presented as follows.

Individual Level

G1) Impasse rate has a positive and linear dependence on goal difficulty.

G2) Agents with specific goals will have higher individual utilities than those with non-specific goals.

G3) Individual utility has a positive and linear dependence on goal difficulty.

In the literature we surveyed, only Huber & Neale (1986) studied the effect of goal difficulty on dyadic performance. They proposed that the goal difficulty of the negotiators would have a positive and additive effect on dyadic performance. When negotiators goals were not sufficiently high, they made concession quickly. When their goals were difficult, they focused on their individual gains rather than on joint profits. Thus, when negotiator goals were homogeneous, highest joint profits would be obtained when both negotiators had moderate goals instead of difficult goals. When both negotiators had difficult goals, the goals were too difficult and too specific to energize performance. However, when both negotiators had moderate goals, high joint profit was achieved because there was still room for integrative agreements (Huber & Neale, 1986). On the other hand, when negotiator goals were disparate, highest joint profits were obtained when one negotiator had a moderate goal while the other had a difficult goal. Under this condition, there was still room for conflict resolution so that highly integrative behavior was observed (Huber & Neale, 1986). All of these findings in human negotiation will be tested in electronic setting in the dyadic level.

Dyadic Level

G4) The assigned goals of both agents will have a positive and additive effect on joint utility.

G5) When agent goals are homogeneous, joint utility will be the highest when both agents have moderate goals.

G6) When agent goals are disparate, joint utility will be the highest when one agent has a difficult goal and the other has a moderate goal.

Pruitt (1981) stated that negotiators with a high level of goal difficulty require a longer time to reach agreements. In order to verify this, the following hypothesis on the relationship between time to reach agreement and goal difficulty will be tested in electronic negotiation.

G7) The assigned goals of both agents will have a positive and additive effect on the time to reach agreements.

Huber & Neale (1986) also investigated the relative effects of goals, by examining the relative profitability of negotiators assigned disparate goals. Since goal difficulty was believed to have a positive and linear relationship with the individual performance of a negotiator, they stated that in disparate goal conditions, negotiators with more difficult goals were more profitable than negotiators with less difficult goals. This hypothesis will be also validated in the electronic negotiation.

Relative Effects of Goals

G8) When goals are disparate between two negotiation agents, the agent with the more difficult goal will achieve higher individual utility than the one with the less difficult goal.

7.5 Joint Effect of Goal Setting and Learning

The joint effect of goal setting and learning has not been studied so far in both human and electronic negotiations. Deveaux *et al.* (2001) found that agents achieved higher performance when they adapted their negotiation strategy to the behavior of their opponents. Thus, higher individual and joint utilities are expected when agents learn.

When the level of goal difficulty increases, the potential for conflict resolution and the room for integrative behavior decrease. Information obtained from learning becomes less helpful. Since agents with more difficult goals already completed fewer number of agreements (i.e., higher impasse rate), there is not as much room for such agents to further adjust their negotiation tactics, unlike agents with less difficult goals. Thus, learning is believed to have higher effect on agents with non-specific and easy goals than on agents with moderate and difficult goals. This is contrary to the observation related to situational power, in which more powerful agents experience greater increases in individual utilities from learning.

Although information from learning might improve an agents' performance, the agent still needs to ensure that each agreement fulfills its goal. Thus, even though learning-induced performance improvement of agents with easy goals might be higher than that with difficult goals, agents with easy goals are still expected to have lower individual performance than those with more difficult goals. Based on the discussion, the following hypotheses are proposed and will be tested on both individual and dyadic levels.

Individual Level

GL1) Agents with learning will have higher individual utilities than those without learning.

GL2) Learning will provide a higher increase in individual utility for agents with less difficult goals, as compared to agents with more difficult goals.

GL3) When agents learn, the agents with more difficult goals will achieve higher individual utilities than those with less difficult goals.

Dyadic Level

GL4) Joint utility with learning will be higher than that in negotiation without learning.

7.6 Experimental Design

7.6.1 Negotiation Scenario

The negotiation in our experiment deals with a supply chain integration between a supplier and a buyer. Ghodsypour and O'Brien (1998) have identified several supplier-selection criteria. Among those criteria, we selected the criteria capability, flexibility, response rate, and defect rate as the non-negotiating issues, and price and lead-time as the negotiating issues. In our experiment, all selling agents provided the same service with undifferentiable properties. The possible properties of the chosen non-negotiating criteria and the values on these criteria for the selling agents, are given in Table 7.2 below.

Table 7.2. Properties of the Service provided by the Selling Agents.

Non-Negotiating Criteria	Possible Properties	Property of Selling Service
Capability	Low, High	High
Flexibility	Low, Moderate, High	Low
Response Rate	Slow, Fast	Fast
Defect Rate	0%, 1%, 2%, 3%, 4%, 5%	0%

Price was the most important negotiating issue for the buying agents while lead time was the most important negotiating issue for the selling agents. In each negotiation, the buying and the selling agents negotiated the price and the lead time of the service, until a compromise was made or either agent terminated the negotiation because its deadline was reached. Each agent was allowed to complete as many transactions as possible, but the maximum number of turns for every agent in every negotiation was 20. In order to keep consistency between human and electronic negotiations, the negotiating ranges of the price and the lead time for both the buying and the selling agents were fixed and are given as

Price: $P_{min} = 100$, $P_{max} = 300$

Lead-time: $LT_{min} = 10$, $LT_{max} = 30$

The values of K were fixed at 0.1 for both the price and the lead time of the agents. In human negotiation, different negotiators had different concession behaviors. We achieved this in electronic negotiation by varying the values of β. The values of β were randomly generated within a range from 0.1 to 7.5. Each agent had a different value of β, but for a given agent, the values of β were the same for both the price and the lead time. In addition, the *filtering thresholds* of all buying agents were fixed at 0.1.

7.6.2 Design for Perceived Power and Learning

The objective of this experiment is to test the hypotheses P1 to P4 as well as PL1 to PL5. The experimental design was similar to that proposed by Pinkley *et al.* (1994). In our experiment, three levels of perceived power, low, equal and high, were proposed.

According to earlier studies, power can be manipulated in many different forms. One way to manipulate power is the incorporation of a negotiator's BATNA. Negotiators are given alternatives with certain levels of utilities when they negotiate. Negotiators with high power have BATNA with high utilities. They might choose to withdraw from negotiations unless the utilities of the negotiated agreements are higher than that of their BATNA (Pinkley *et al.*, 1994). Another way to manipulate power is to vary the number of buyers and sellers in the market (McAlister *et al.*, 1986; Dwyer & Walker, 1981). When there are fewer buyers than sellers in the market, buyers have higher power than sellers have, because sellers need to compete with other sellers to make successful agreements while buyers have more options to choose. McAlister *et al.* (1986) has further manipulated the power economically by setting the maximum number of transactions that buyers and sellers are allowed to complete. The smaller the number of transactions that a negotiator can make, the higher the power the negotiator has. However, this power manipulation method has been questioned, because power is manipulated as a function of individual level behavior instead of market level behavior (Pinkley *et al.*, 1994).

Instead of using perceived final utilities as Pinkley *et al.* (1994) did, in our experiment, an agent's perception of power is represented by the preference of the market for the attributes of the product. If an agent perceives that the market prefers those attributes, the buyer (seller) has low (high) perceived power.

By using such a manipulation of power, a 3 x 3 factorial experimental design was established, with three levels of perceived power of both buying and selling agents as two factors. In each of the resulting nine treatment combinations, there were 150 replications, and in each experimental replication, there were five buying agents and five selling agents, resulting a total of 25 dyads. An overview of the 3 x 3 experimental design is presented in Table 7.3a. Parameters that were held constant in all treatments are given at the bottom of the table.

The preference levels of the buying and the selling agents are presented in Table 7.3b. Since price was the most important issue for the buying agents, the preference level of price of all buying agents was set to be seven times more important than that of lead time and that of other non-negotiating issues. The lead time was set to be equally important as the other non-negotiating issues. In addition, the preference level between any two of the attributes from the same non-negotiating issue was equally important. On the other hand, since lead time was the most important issue for the selling agents, the preference level of lead time of all selling agents was seven times more important than that of price.

The perceived market preferences corresponding to the three levels of perceived power are given in Table 7.3c. With SP 1 (column four), it is perceived that buyers in the market strongly prefer a service with high capability, low flexibility, fast response rate, and 0% defect. Since these are the attributes of the selling service as well (column two), buyers in the market are perceived to have a higher preference for the selling service, as compared to the neutral preferences in the last column. Thus, buying agents whose perceived market preferences are given by SP 1 have low perceived power, while selling agents with a perception of SP 1 have high perceived power.

On the other hand, with SP 3 in column six, it is perceived that buyers in the market do not prefer service with high capability, low flexibility, fast response rate, and 0% defect. A comparison with the neutral preferences indicates that buyers in the market are perceived to have a lower preference on the selling service. As a result, buying agents with SP 3 as the perceived market preferences have high perceived power, while selling agents which perceive SP 3 have low perceived power.

When we compare SP 2 in column five with the neutral preference, the perceived market preferences and the preferences of the neutral buyer are the same. Thus, both buying agents and selling agents are said to have equal power when they perceive the market preferences as given by SP 2.

The experimental design of investigating the joint effect of power and learning was the same as discussed above except for one thing. In all experimental treatments, both buying and selling agents were able to learn.

Table 7.3. Complete Experimental Design for Perceived Power and Learning.

Table 7.3a. Overview of the 3 x 3 Experimental Design for Perceived Power and Learning.*

Buying Agent's Perceived Power	Selling Agent's Perceived Power		
	High (SP 1)	**Equal (SP 2)**	**Low (SP 3)**
Low (SP 1)	Low / High	Low / Equal	Low / Low
Equal (SP 2)	Equal / High	Equal / Eqaul	Equal / Low
High (SP 3)	High / High	High / Eqaul	High / Low

*Parameters common to all treatments are as follows:

$\beta \in$ [0.1, 7.5]
T_{max} = 20
K = 0.1 for both price and lead-time

P_{min} = 100, P_{max} = 300
LT_{min} = 10, LT_{min} = 30

Max. % of Change of β = 50%
Filtering Threshold = 0.1

Table 7.3b. Individual Preference Levels among the Criteria.

Criteria to be Compared	Preference Level	
	Buying Agent	**Selling Agent**
Price vs. Lead Time	7	1/7
Price vs. Any One of the Non-Negotiating Issues	7	--
Lead Time vs. Any One of the Non-Negotiating Issues	1	--
Between Any Two of the Non-Negotiating Issues	1	--
Between Any Two of the Attributes from the Same Non-Negotiating Issue	1	--

Table 7.3c. Perceived Market Preferences among the Properties of each Criterion.

Criteria	Properties of Selling Service	Properties to be Compared	Perceived Market Preference			Neutral Preference
			SP1	**SP2**	**SP3**	
Capability	High	High vs. Low	7	1	1/7	1
Flexibility	Low	Low vs. Moderate Low vs. High	7	1	1/7	1
Response Rate	Fast	Fast vs. Slow	7	1	1/7	1
Defect	0%	0% vs. Others	7	1	1/7	1

7.6.3 Design for Goal Setting and Learning on Electronic-based Negotiation

This experiment was designed based on the setting proposed by Huber & Neale (1987). Four levels of goal difficulty were proposed by varying the *final thresholds* that agents had to achieve in their negotiations. They were do-your-best (non-specific), easy, moderate and difficult goals. A 4 x 4 factorial experimental design was established with buying agent's goal and selling agent's goal as the two factors. In each of the 16 experimental treatments, there were 150 experimental replications, and in each experimental replication, there were five buying agents and five selling agents, resulting in a total of 25 dyads. An overview of the 4 x 4 experimental design is presented in Table 7.4a. Parameters that were held constant to all treatments are given at the bottom of the table.

The preference levels of the buying and the selling agents are presented in Table 7.4b. Since price was the most important issue for the buying agents, the preference level of price of all buying agents was set to be seven times more important than that of the non-negotiating issues. For the buying agents, lead-time was the least important issue. Thus, the preference levels between lead time and price as well as between lead time and other non-negotiating issues were 1/7. In addition, the preference level between any two of the attributes from the same non-negotiating issue was equally important. For the selling agents, the preference level of lead time was seven times more important than that of price.

To make our goal-setting manipulation similar to that of Huber & Neale (1986), we set the goal difficulty with reference to the joint utility for a fully integrative agreement. To explain further, we first observe that when the final price is $100 and the final lead-time is 10 days, buying agents achieve their highest utility (0.9056) and selling agents achieve their lowest utility (0). When the final price is $300 and the final lead-time is 30 days, selling agents achieve their highest utility (1) and buying agents achieve their lowest utility (0). Fully integrative agreement is reached when the final price is $100 and the final lead-time is 30 days. Under this condition, the joint utility reaches 1.753. Buying agents achieve a utility of 0.878 while selling agents achieve a utility of 0.875. According to this information, the level of goal difficulty and the corresponding range of the *final threshold* are presented in Table 7.4c. The *final threshold* of each agent was randomly generated within the defined range.

The experimental design of investigating the joint effect of goal setting and learning was the same as discussed above except one thing. In all experimental treatments, both buying and selling agents were able to learn.

Table 7.4. Complete Experimental Design for Goal Setting and Learning.

Table 7.4a. Overview of the 4 x 4 Experimental Design for Goal Setting and Learning.*

Buying Agent's Goal	**Selling Agent's Goal**			
	No Goal	**Easy**	**Moderate**	**Difficult**
No Goal	No Goal / No Goal	Easy / No Goal	Moderate / No Goal	Difficult / No Goal
Easy	No Goal / Easy	Easy / Easy	Moderate / Easy	Difficult / Easy
Moderate	No Goal / Moderate	Easy / Moderate	Moderate / Moderate	Difficult / Moderate
Difficult	No Goal / Difficult	Easy / Difficult	Moderate / Difficult	Difficult / Difficult

*Parameters common to all treatments are as follows:

$\beta \in [0.1, 7.5]$	$P_{min} = 100$, $P_{max} = 300$	$LT_{min} = 10$, $LT_{min} = 30$
$T_{max} = 20$	Filtering Threshold = 0.1	$K = 0.1$ for both price and lead-time

Table 7.4b. The Individual Preference Levels of both the Buying and the Selling Agents

Criteria to be Compared	**Importance Level**	
	Buying Agent	**Selling Agent**
Price vs. Lead Time	7	1/7
Price vs. Any One of the Non-Negotiating Issues	7	--
Lead Time vs. Any One of the Non-Negotiating Issues	1/7	--
Between Any Two of the Non-Negotiating Issues	1	--
Between Any Two of the Attributes from the Same Non-Negotiating Issue	1	--

Table 7.4c. Final Threshold values for the Goal-Setting Manipulation

Goal Difficulty	**Range of Final Threshold**
No Goal	0
Easy	[0.650, 0.700]
Moderate	[0.750, 0.800]
Difficult	[0.850, 0.885]

7.6.4 Computation of Performance Measures

This section describes the exact manner by which the performance measures are computed in the experimental designs given in §§7.6.2 and 7.6.3. We define

N	Number of replications in each experimental treatment
k	Replication index of the experimental treatment (k^{th} replication)
N_B^k	Number of buying agents in the k^{th} replication of the experimental treatment
N_S^k	Number of selling agents in the k^{th} replication of the experimental treatment
N_{NS}^k	Number of negotiations that has not been started in the k^{th} replication of the experimental treatment, because of non-fulfillment of the *filtering threshold*
N_C^k	Number of completed negotiations in the k^{th} replication of the experimental treatment
i	Buying agent number in each replication of the experimental treatment (i^{th} buying agent)
j	Selling agent number in each replication of the experimental treatment (j^{th} selling agent)
P	Set of treatments for the power-setting experiment P = {Low, Equal, High}
PL	Set of treatments for the power-setting-with-learning experiment PL = {Low with Learning, Equal with Learning, High with Learning}
G	Set of treatments for the goal-setting experiment G = {Non-Specific, Easy, Moderate, Difficult}
GL	Set of treatments for the goal-setting-with-learning experiment GL = {Non-Specific and Learning, Easy and Learning, Moderate and Learning, Difficult and Learning}

The performance measure of each experimental treatment is denoted as PM(b, s), where PM stands for one of the performance measures given in Table 7.1, and b and s are the experimental conditions of the buying and the selling agents, respectively, as given below.

$b, s \in P$ for power setting experiment, or

$b, s \in PL$ for power setting with learning experiment, or

$b, s \in G$ for goal setting experiment, or

$b, s \in GL$ for goal setting with learning experiment.

The equation for the *average* performance measure of each experimental treatment is obtained as

$$PM(b,s) = \frac{\sum_{k=1}^{N} PM(b,s)^k}{N} \tag{7.3}$$

where $PM(b,s)^k$ represents the performance measure of the k^{th} replication of the experimental treatment and is obtained by

$$PM(b,s)^k = \frac{\sum_{i=1}^{N_B^k} \sum_{j=1}^{N_S^k} PM(b,s)_{ij}^k}{N_B^k N_S^k - N_{NS}^k} \tag{7.4}$$

where $PM(b,s)_{ij}^k$ is the performance measure of the dyad (i, j) at the k^{th} replication, and is obtained as given in §7.1. To avoid biasing the experimental result, we take account of impasse rate in utility computation (Tripp & Sondak, 1992). When agents break off a negotiation, they receive zero utilities, as suggested in the literature (Hamner and Harnett, 1974; Pinkley *et al.*, 1994).

Equation (7.4) is not applicable to the measure of impasse rate, because impasse rate is not measured for each dyad. Instead, the impasse rate of the k^{th} replication of the experimental treatment is obtained by

$$IR(b,s)^k = \frac{N_B^k N_S^k - N_{NS}^k - N_C^k}{N_B^k N_S^k - N_{NS}^k} \tag{7.5}$$

Since our experiments consist of multiple issues, searching for the Nash solution is a complex nonlinear problem. To simplify this, we estimate the Nash solution by enumerating over the feasible ranges of the price and the lead time for each negotiation, and compute the joint utility for each possible combination. We then take the agreement with the highest joint utility as the approximate Nash solution for that dyad.

8 Results for Power Setting Experiment

In §7.6.2, we discussed the experimental design for testing the effects of situational power and learning on electronic negotiation. In this chapter, we conduct the experiment, and present the associated results and discussions.

For the nine-treatment experiment without learning, the average values of the buying agent's utility, the selling agent's utility, the joint utility, the number of turns and the distance to Nash are shown in Table 8.1. When both buying agents and selling agents agree on their perceptions of the market's preferences, the average number of turns to reach agreements is around seven for our experimental setting. Buying agents achieve the highest individual utility (M = 0.49), when both parties agree that buyers in the market do not like the product. Selling agents achieve the highest individual utility (M = 0.65), when both parties agree that buyers in the market like the product. In addition, when the perception of power of one party increases while that of the other party remains unchanged, individual utility increases for the party with increasing power.

The average performance measures for the nine experimental treatments with learning are given in Table 8.2. A similar pattern is obtained when compared with Table 8.1. Again, buying agents achieve the highest individual utility (M = 0.50), when both parties agree that buyers in the market do not like the product, while selling agents achieve the highest individual utility (M = 0.71), when both parties agree that buyers in the market like the product. Similarly, when both agents learn, the higher the power a party perceives, the higher the individual utility the party achieves, when the perceived power of the other party remains unchanged. By comparing the performance measures between the learning and no-learning situations, we see that agents with learning achieve higher individual utilities, higher joint utilities and fairer agreements. With these observations, we now proceed to test the hypotheses given in §§7.2 and 7.3.

Table 8.1. Results for Various Treatment Combinations of Perceived Power without Learning.

Buying Agent's Power (b)	Performance Measure	Selling Agent's Power (*s*)					
		High Power (SP 1)		Equal Power (SP 2)		Low Power (SP 3)	
		Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Low Power (SP 1)	BU(*b*, *s*)	0.36	0.03	0.41	0.03	0.44	0.03
	SU(*b*, *s*)	0.65	0.07	0.54	0.07	0.46	0.07
	JU(*b*, *s*)	1.01	0.04	0.94	0.04	0.90	0.04
	NT(*b*, *s*)	7.26	1.28	5.92	1.10	5.14	0.99
	DN(*b*, *s*)	0.25	0.07	0.35	0.07	0.42	0.07
Equal Power (SP 2)	BU(*b*, *s*)	0.39	0.03	0.43	0.03	0.46	0.03
	SU(*b*, *s*)	0.59	0.07	0.47	0.07	0.40	0.07
	JU(*b*, *s*)	0.98	0.05	0.90	0.05	0.86	0.04
	NT(*b*, *s*)	8.58	1.35	6.90	1.21	5.90	1.15
	DN(*b*, *s*)	0.30	0.07	0.41	0.07	0.48	0.07
High Power (SP 3)	BU(*b*, *s*)	0.42	0.03	0.47	0.03	0.49	0.02
	SU(*b*, *s*)	0.49	0.07	0.38	0.07	0.32	0.06
	JU(*b*, *s*)	0.92	0.05	0.85	0.04	0.81	0.04
	NT(*b*, *s*)	10.52	1.43	8.29	1.39	7.00	1.34
	DN(*b*, *s*)	0.40	0.06	0.51	0.06	0.56	0.06

Table 8.2. Results for Various Treatment Combinations of Perceived Power with Learning.

Buying Agent's Power (b)	Performance Measure	Selling Agent's Power (s)					
		High Power (SP 1)		Equal Power (SP 2)		Low Power (SP 3)	
		Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Low Power (SP 1)	BU(*b*, *s*)	0.40	0.02	0.43	0.02	0.45	0.43
	SU(*b*, *s*)	0.71	0.06	0.58	0.07	0.50	0.58
	JU(*b*, *s*)	1.11	0.04	1.01	0.05	0.94	1.01
	NT(*b*, *s*)	12.01	2.04	8.92	2.25	7.27	8.92
	DN(*b*, *s*)	0.23	0.04	0.32	0.05	0.40	0.32
Equal Power (SP 2)	BU(*b*, *s*)	0.43	0.02	0.45	0.02	0.47	0.43
	SU(*b*, *s*)	0.69	0.06	0.55	0.08	0.46	0.69
	JU(*b*, *s*)	1.12	0.05	1.01	0.06	0.93	1.12
	NT(*b*, *s*)	13.81	1.72	10.55	2.21	8.59	13.81
	DN(*b*, *s*)	0.23	0.05	0.34	0.06	0.43	0.23
High Power (SP 3)	BU(*b*, *s*)	0.45	0.02	0.48	0.02	0.50	0.45
	SU(*b*, *s*)	0.65	0.08	0.51	0.08	0.41	0.65
	JU(*b*, *s*)	1.11	0.06	1.00	0.07	0.91	1.11
	NT(*b*, *s*)	15.57	1.39	12.67	2.05	10.46	15.57
	DN(*b*, *s*)	0.25	0.07	0.38	0.07	0.47	0.25

Note: *BU stands for buyer's IU, and SU stands for seller's IU*
Each performance measure in the above tables is obtained using (7.3) in §7.6.4.

8.1 Effect of Perceived Power

Individual Performance

Hypothesis P1. (Higher-Power Agent's IU > Lower-Power Agent's IU) We test the effect of perceived power on average individual utility per transaction by the regression equation:

$$\text{AVG\_UTILITY} = \lambda_0 + \lambda_1\text{ROLE} + \lambda_2\text{POWER} + \varepsilon$$

where AVG_UTILITY is the k^{th} replication's average utility, $\text{IU}(b, s)^k$, for b, $s \in P$, ROLE is the buyer (seller) if IU stands for the buyer's (seller's) IU, POWER is the perceived power given by b (s) if IU stands for the buyer's (seller's) IU, λ_0, λ_1 and λ_2 are the regression coefficients, and ε is the error term that is normally distributed with the mean. POWER is a dummy variable coded as 1 for low, 2 for equal, and 3 for high. In all, there are 1350 values of buyer's IUs (= 150 replications/treatment x 9 treatments) and 1350 values of seller's IUs, and by using these values, we find the regression equation as:

$$\text{AVG\_UTILITY} = 0.311 + \underset{(p < 0.0005)}{0.04767\text{ROLE}} + \underset{(p < 0.0005)}{0.05904\text{POWER}} + \varepsilon$$

The equation is significant ($p < 0.0005$, $R^2 = 0.329$)[2]. As the coefficient of the POWER variable is positive and significant, the average individual utility per transaction increases as the level of power increases. The ROLE variable is also significant, because buying and selling agents have different payoff tables in this experiment.

We further test this hypothesis by comparing the individual utilities among agents with three levels of power. These treatments corresponding to different power settings are:

Low-Power Agents:	$\{(b, s) \mid b = \text{Low}, s \in P\}$	for buyer	(8.1a)
	$\{(b, s) \mid b \in P, s = \text{Low}\}$	for seller	(8.1b)

[2] Discussion for a low value of R^2 in this hypothesis and the other hypotheses is presented in §8.3

Equal-Power Agents:	$\{(b, s) \mid b = \text{Equal}, s \in P\}$	for buyer	(8.2a)
	$\{(b, s) \mid b \in P, s = \text{Equal}\}$	for seller	(8.2b)
High-Power Agents:	$\{(b, s) \mid b = \text{High}, s \in P\}$	for buyer	(8.3a)
	$\{(b, s) \mid b \in P, s = \text{High}\}$	for seller	(8.3b)

where each tuple (b, s) in the above equations consists of 150 replications, and represents an experimental treatment in Table 8.1.

There are 450 values of IUs (= 150 replications x 3 treatments of power level) for each of (8.1a) through (8.3b). Thus, each power level in the ANOVA test contains 900 values. We observe significant differences among the average individual utilities corresponding to the three power levels ($p < 0.0005$). The results are presented in Tables 8.3. Post-hoc results in Table 8.4 further show that agents with high power ($M = 0.5169$) achieve significantly higher individual utilities than agents with equal power ($M = 0.4439$, $p < 0.0005$), while agents with equal power achieve significantly higher individual utilities than agents with low power ($M = 0.3988$, $p < 0.0005$). Overall, this hypothesis is supported.

Table 8.3. Individual Utilities for the Power-and-No-Learning Environment.

Perceived Power	Individual Utility
Low	0.3988
Equal	0.4439
High	0.5169

Note: Individual utilities among three levels of power are found to be significant (p < 0.0005).

Table 8.4. Post-hoc Results for Individual Utilities for the Power-and-No-Learning Environment.

Comparisons	p-value
Low vs. Equal	< 0.0005
Low vs. High	< 0.0005
Equal vs. High	< 0.0005

Dyadic Performance

Hypothesis P2. (Balanced Power Agreement's JU > Imbalanced Power Agreement's JU)
Balanced power agreements consist of those treatments where both agents perceive the same

level of power, while in an imbalanced power agreement, one party perceives higher power than the other party does. They are given by:

Balanced Power Agreements:	$\{(b, s) \mid b = s \text{ for } b, s \in P\}$	(8.4)
Imbalanced Power Agreements:	$\{(b, s) \mid b \neq s \text{ for } b, s \in P\}$	(8.5)

We further divide (8.5) into two groups as given by

Imbalanced Power Favoring Buyers:	$\{(b, s) \mid b > s \text{ for } b, s \in P\}$	(8.5a)
Imbalanced Power Favoring Sellers:	$\{(b, s) \mid b < s \text{ for } b, s \in P\}$	(8.5b)

With 150 experimental replications in each tuple (b, s), there are 450 values of JUs in (8.4) (= 150 x 3 balanced power treatments) and 900 values (= 150 x 6 imbalanced power treatments) in (8.5); (8.5a) and (8.5b) consist of 450 values.

Table 8.5 shows the JUs for the treatments with imbalanced power agreements. Results from a t-test indicates that there is no significant difference between (8.4) and (8.5). However, when we only compare the balanced power agreements with the imbalanced power agreements in (8.5b), we observe that the joint utilities in the former one are significantly lower ($p < 0.0005$). On the other hand, the balanced power agreements have significantly higher joint utilities than those agreements with imbalanced power favoring buyers ($p < 0.0005$). Since the decrease in the joint utilities from the agreements with imbalanced power favoring sellers is cancelled out by the increase in the joint utilities from the agreements with imbalanced power favoring buyers, there is no significant difference in the joint utilities between the balanced and the imbalanced power agreements. Hence, this proposition is partially supported.

Table 8.5. Joint Utilities of Imbalanced Power Agreements.

Balanced versus	**Joint Utility**	**t-test p-value**
Imbalanced	0.9067	0.834
Imbalanced Favoring Sellers	0.9756	< 0.0005
Imbalanced Favoring Buyers	0.8377	< 0.0005

Note: Mean JU for Balanced Power Agreements is 0.9058.

Hypothesis P3. (Self and Competitor's Power have a Positive and Additive Effect on NT) We test the effect of perceived power on the number of turns by the regression equation:

$$\text{AVG\_TURNS} = \lambda_0 + \lambda_1\text{BPOWER} + \lambda_2\text{SPOWER} + \lambda_3\text{POWER\_INT} + \varepsilon$$

where AVG_TURNS is the k^{th} replication's average number of turns, $NT(b, s)^k$, for b, s $\in$ P, BPOWER and SPOWER are dummy variables for the buyer's and seller's perceived powers (b and s), and POWER_INT is the interaction between the two dummy variables. This means that there are 1350 values of NTs (= 150 x 9 treatments). We find the regression equation as:

$$\text{AVG\_TURNS} = 3.420 + \underset{(p < 0.0005)}{0.543\text{BPOWER}} + \underset{(p < 0.0005)}{0.681\text{SPOWER}} + \underset{(p < 0.0005)}{0.353\text{POWER\_INT}} + \varepsilon$$

The equation is significant ($p < 0.0005$, $R^2 = 0.596$). The coefficients of BPOWER and SPOWER dummy variables are positive and significant. This hypothesis is supported.

Hypothesis P4. (Balanced Power Agreement's DN < Imbalanced Power Agreement's DN) To test this proposition, we compare the distance to Nash for (8.4) with that for (8.5). The smaller the distance to Nash, the fairer the agreement. Among nine experimental treatments, the three treatments consisting of balanced power agreements yield an average distance to Nash of 0.4139. Table 8.6 shows the distances to Nash for treatments consisting of imbalanced power agreements. Results from t-tests indicate that there is no significant difference in the distances to Nash between the two treatments.

However, when we compare the treatments consisting of balanced power agreements with the treatments where the selling agents perceive higher power, we observe that the distances to Nash in the former one are significantly higher ($p < 0.0005$). On the other hand, the treatments consisting of balanced power agreements have significantly lower distances to Nash than those treatments where the buying agents perceive higher power ($p < 0.0005$). That is, in the overall imbalanced power treatments, an increase in the distance to Nash for treatments favoring sellers cancels out the decrease in the distance for treatments favoring buyers.

In summary, although there appears to have no difference between balanced and imbalanced power treatments, balanced power agreements are fairer than imbalanced ones favoring buyers. Thus, this proposition is only partially supported.

Table 8.6. Distances to Nash of Imbalanced Power Agreements.

Balanced versus	**Distance to Nash**	**t-test p-value**
Imbalanced	0.4084	0.415
Imbalanced Favoring Sellers	0.3001	< 0.0005
Imbalanced Favoring Buyers	0.5168	< 0.0005

Note: Mean Distance to Nash for Balanced Power agreements is 0.4139.

8.2 Joint Effect of Perceived Power and Learning

Individual Performance

Hypothesis PL1. (IU of Agents with Learning > IU of Agents without Learning) To test this proposition, we employ a t-test and regression analysis, and the results from both tests support the hypothesis. First, we compare the individual utilities between agents with and without learning using a t-test. These agents are given by:

Agents without Learning: $\{(b, s) \mid b, s \in P\}$ (8.6)

Agents with Learning: $\{(b, s) \mid b, s \in PL\}$ (8.7)

There are 2700 values of IUs for each of (8.6) and (8.7). Results indicate that agents with learning achieve significantly higher individual utilities (M = 0.5076) than those without learning (M = 0.4532, p < 0.0005).

We further test the effect of learning on average individual utility per transaction by the regression equation:

$$\text{AVG\_UTILITY} = \lambda_0 + \lambda_1\text{ROLE} + \lambda_2\text{POWER} + \lambda_3\text{LEARNING} + \lambda_4\text{POWER\_LEARNING\_INT} + \varepsilon$$

where AVG_UTILITY is the k^{th} replication's average individual utility, $IU(b, s)^k$, for b, s $\in P \cup PL$, LEARNING indicates whether the agent learns or not, POWER_LEARNING_INT is the interaction between power and learning, and other dummy variables are defined in the same way as before. We find the regression equation as:

$$\text{AVG\_UTILITY} = 0.284 + \underset{(p < 0.0005)}{0.08050\text{ROLE}} + \underset{(p < 0.0005)}{0.04818\text{POWER}} + \underset{(p < 0.0005)}{0.03265\text{LEARNING}} + \underset{(p < 0.0005)}{0.01086\text{POWER\_LEARNING\_INT}} + \varepsilon$$

This equation is significant ($p < 0.0005$, $R^2 = 0.476$), and the coefficient of the LEARNING dummy variable is positive and significant as well. In sum, this proposition is supported.

Hypothesis PL2. (Higher-Power Agents have Higher Increase in Utility from Learning) To test this proposition, we first compare the individual utilities between agents with and without learning when these agents have different levels of perceived power. Agents with learning and different levels of perceived power are given by:

Low-Power Agents with Learning:	$\{(b, s) \mid b = \text{Low with Learning}, s \in PL\}$	for buyer	(8.8a)
	$\{(b, s) \mid b \in PL, s = \text{Low with Learning}\}$	for seller	(8.8b)
Equal-Power Agents with Learning:	$\{(b, s) \mid b = \text{Equal with Learning}, s \in PL\}$	for buyer	(8.9a)
	$\{(b, s) \mid b \in PL, s = \text{Equal with Learning}\}$	for seller	(8.9b)
High-Power Agents with Learning:	$\{(b, s) \mid b = \text{High with Learning}, s \in PL\}$	for buyer	(8.10a)
	$\{(b, s) \mid b \in PL, s = \text{High with Learning}\}$	for seller	(8.10b)

where each tuple (b, s) in the above equations represents the experimental treatment in Table 8.2, and consists of 150 experimental replications.

First, we investigate if learning can improve individual utilities for agents with different levels of perceived power. We conducted t-tests to compare the individual utilities between (8.1a ∪ 8.1b) and (8.8a ∪ 8.8b), (8.2a ∪ 8.2b) and (8.9a ∪ 8.9b), and (8.3a ∪ 8.3b) and (8.10a ∪ 8.10b). The results are presented in Table 8.7. We observe that agents with low power, equal power, and high power all achieve significantly higher individual utilities when they learn.

Table 8.7. Effect of Learning on Individual Utilities at Different Levels of Perceived Power.

Perceived Power	Individual Utility		Average Increase	t-test p-value
	Without Learning	With Learning		
Low	0.3988	0.4417	0.0429	< 0.0005
Equal	0.4439	0.4996	0.0557	0.001
High	0.5169	0.5815	0.0646	< 0.0005

Next, we test if the *increases* in utilities are different among the three power levels. We conducted a 2 x 3 two-way ANOVA test that consists of agents with and without learning at three different levels of perceived power. We compare the *increases* in individual utilities for the three levels of power. The results are presented in Figure 8.1, and indicate an interaction effect between perceived power and learning.

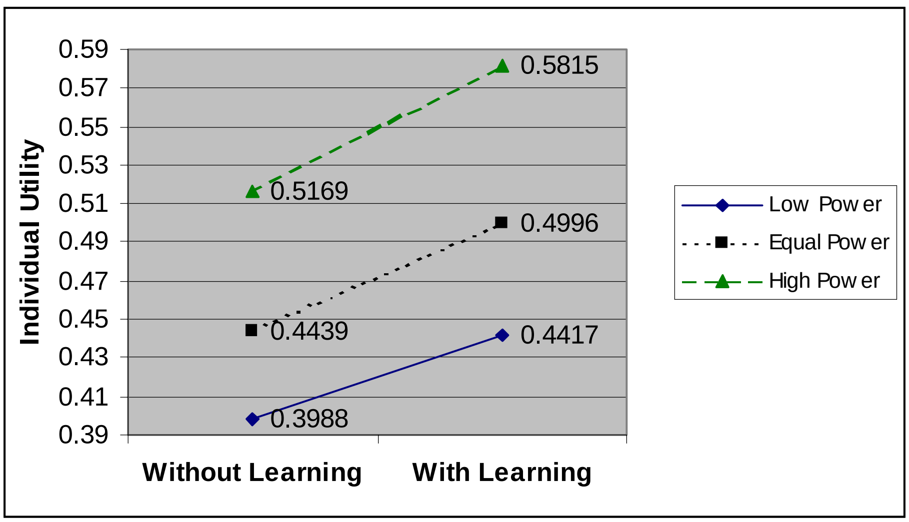

Figure 8.1. Interaction Effect of Learning and Power Perception on Individual Utility.

That is, when agents learn, agents with high power achieve higher increases in individual utilities (M = 0.0646) than agents with equal power (M = 0.0557), while agents with equal power achieve higher increases in individual utilities than those with low power (M = 0.0429). In sum, this hypothesis is supported.

Hypothesis PL3. (When Agents Learn, Higher-Power Agent's IU > Lower-Power Agent's IU) To test this hypothesis, we compare the individual utilities among low-power agents with learning (8.8a ∪ 8.8b), equal-power agents with learning (8.9a ∪ 8.9b), and high-power agents with learning (8.10a ∪ 8.10b). A one-way ANOVA test indicates significant difference among these agents ($p < 0.0005$), and the results are presented in Table 8.8. Post-hoc results in Table 8.9 further show that when agents learn, the individual utilities of high-power agents (M = 0.5815) are significantly higher than those of equal-power agents (M = 0.4995, $p < 0.0005$), and those of low-power agents (M = 0.4417, $p < 0.0005$). When equal-power agents are compared with low-power agents, equal-power agents achieve significantly higher individual utilities. Overall, this hypothesis is supported.

Table 8.8. Individual Utilities for Three Levels of Power in Learning Environment.

Perceived Power	Individual Utility
Low	0.4417
Equal	0.4995
High	0.5815

Note: Individual utilities among three levels of power are found to be significant (p < 0.0005).

Table 8.9. Post-hoc Results for Individual Utilities at Three Levels of Power in Learning Environment.

Comparisons	p-value
Low vs. Equal	< 0.0005
Low vs. High	< 0.0005
Equal vs. High	< 0.0005

Dyadic Performance

Hypothesis PL4. (JU of Agreements with Learning > JU of Agreements without Learning, and DN of Agreements with Learning < DN of Agreements without Learning) When agents learn, balanced power agreements and imbalanced power agreements are given by:

Balanced Power Treatments with Learning: $\{(b, s) \mid b = s \text{ for } b, s \in PL\}$ (8.11)

Imbalanced Power Treatments with Learning: $\{(b, s) \mid b \neq s \text{ for } b, s \in PL\}$ (8.12)

Since there are 150 experimental replications for each tuple in the above equations, there are 450 values of JUs and DNs (= 150 x 3 balanced power treatments) for (8.11), and 900 values of JUs and DNs (= 150 x 6 imbalanced power treatments) for (8.12). To test this hypothesis, we compare the joint utilities and the distances to Nash for (8.4) vs. (8.11) and for (8.5) vs. (8.12).

We observe that balanced power agreements with learning have significantly higher joint utilities (M = 1.0197) and lower distances to Nash (M = 0.3311) than those without learning (M = 0.9058, p < 0.0005 for JU, and M = 0.4139, p < 0.0005 for DN). In imbalanced power conditions, agreements with learning also have significantly higher joint utilities (M = 1.0129) and lower distances to Nash (M = 0.3425) than those without learning (M = 0.9067, p < 0.0005 for JU, and M = 0.4084, p < 0.0005 for DN). On the whole, this proposition is supported.

Table 8.10. Effect of Learning and Power on Joint Utilities and Distances to Nash.

Performance Measure	Power Perception	No Learning	Learning	T-test p-value
Joint Utility	Balanced Imbalanced	0.9058 0.9067	1.0197 1.0129	< 0.0005 < 0.0005
Distance to Nash	Balanced Imbalanced	0.4139 0.4084	0.3311 0.3425	< 0.0005 < 0.0005

8.3 Discussion

8.3.1 Effect of Perceived Power

A summary of the results of the effect of perceived power on electronic negotiation outcome is presented in Table 8.11, wherein the shaded portion indicates hypotheses that are not fully supported. According to our results, negotiation agents with higher (lower) perceived power achieve higher (lower) utilities, but take longer (less) time to reach agreements. This implies that there is a trade-off between performance and time when we integrate situational power in electronic negotiation.

Dwyer & Walker (1981), and McAlister *et al.* (1986) found that the joint utility in an imbalanced power situation is lower than that in a balanced power situation. On the other hand, Pinkley *et* *al.* (1994) proposed that the joint utility is higher in an imbalanced power situation. Our results indicate that there is no significant difference in the joint utilities between balanced and imbalanced power agreements. According to Table 8.8, when selling agents perceive higher power than buying agents do, the joint utilities increase when compared with the joint utilities in balanced power conditions. In another case, when buying agents perceive higher power than selling agents do, the joint utilities decrease when compared with the joint utilities in balanced power conditions. Although our results are not totally consistent with the findings from Dwyer & Walker (1981) and McAlister *et al.* (1986), our results do not support the proposition from Pinkley *et al.* (1994) either.

The discrepancy is caused by the difference in the payoff tables between the buying and the selling agents. To compute utilities, buying agents incorporate the values of both the negotiating and the non-negotiating issues, but selling agents consider the negotiating issues only. In addition, there is only one type of product for buying agents to choose, while selling agents can sell their products to any buying agent. Thus, sellers have an advantage, and hence, perceived power has a greater effect on selling agents than on buying agents. In other words, selling agents with high-perceived power achieve higher increases in utilities than buying agents with high-perceived power. This explains why the joint utilities increase when selling agents have higher perceived power than buying agents do. Because of this, when selling agents have higher perceived power, agreements become more integrative, and the joint utilities approach closer to the integrative level. Thus, the distances to Nash decrease. This explains why agreements become fairer when selling agents have higher perceived power.

Finally, we observe that in some cases, the value of R^2 is our regression equation is comparatively low, especially when we deal with the utility of a negotiation. An explanation for this is that agents have different concession behaviors in our experiment. According to our design, we vary the value of β from 0.1 to 7.5 for each of the three levels of power. Faratin *et al.* (1997) have determined that this variation can affect an agent's negotiation performance.

Table 8.11. Summary of the Effects of Perceived Power on Electronic Negotiation Outcome.

	Hypothesis	Our Results	Literature supporting Hypothesis	Literature not supporting Hypothesis
P1	Agents with higher perceived power will have higher individual utilities than those with lower perceived power.	Supported	Pinkley *et al.* (1994) Dwyer & Walker (1981) McAlister *et al.* (1986)	None
P2	Joint utility in a balanced power condition will be higher than that in an imbalanced power condition.	Partially Supported	Dwyer & Walker (1981) McAlister *et al.* (1986)	Pinkley *et al.* (1994)
P3	The levels of perceived power of both agents will have a positive and additive effect on the time to reach agreements	Supported	Dwyer & Walker (1981)	None
P4	Distance to Nash solution in a balanced power condition will be lower than that in an imbalanced power condition.	Partially Supported	None	None

8.3.2 Joint Effect of Perceived Power and Learning

Table 8.12 presents the summary of the results of the joint effect of perceived power and learning on electronic negotiation outcome. Our heuristic algorithms for learning are found to be effective for agents with low, equal, and high perceived power. Higher individual utilities are achieved when agents learn. Moreover, in both balanced and imbalanced power conditions, agreements become more integrative and fairer when agents carry out the proposed learning algorithms. In general, learning is beneficial for agents, irrespective of their levels of power perception.

We also observe that learning provides a higher increase in individual utility for agents with higher power, when compared to agents with lower power. This implies that learning is more effective for high-power agents than for low-power agents in electronic negotiation.

When learning is implemented, we see that agents with higher power achieve higher utilities than those with lower power. Although we find that learning is effective in improving the individual performance of an agent, situational power remains in effect when agents learn. Thus, we suggest incorporating situational power into agents, even agents learn, and vice versa.

Table 8.12. Summary of the Joint Effect of Perceived Power and Learning on Electronic Negotiation Outcome.

	Hypothesis	Our Results	Literature supporting Hypothesis	Literature not supporting Hypothesis
PL1	Agents with learning will have higher individual utilities than those without learning.	Supported	Zeng & Sycara (1998) Deveaux et al. (2001)	None
PL2	Learning will provide a higher increase in individual utility for agents with higher power, as compared to agents with lower power.	Supported	None	None
PL3	When learning is implemented, agents with higher power will have higher individual utilities than those with low power.	Supported	None	None
PL4	In both balanced and imbalanced power conditions, agreements with learning will have higher joint utilities, and lower distances to Nash than those without learning.	Supported	Zeng & Sycara (1998) supported the proposition on Nash	None

9 Results for Goal Setting Experiment

In Chapter 8, we discussed the effects of situational power and learning on electronic negotiation. In this chapter, with the experimental design in §7.6.3, we conduct the experiment to test the effect of goal constraints on electronic negotiation. Also, we present results and discussion.

For the 16 goal-setting combinations without learning, Table 9.1 shows the average values of the buying agent's utility, the selling agent's utility, the joint utility, the number of turns and the impasse rate. Generally, as the difficulty of an agent's goal increases, a higher individual utility is achieved, but a longer time is required and a higher impasse rate is expected. Joint utility is the highest when selling agents have a difficult goal while buying agents have a moderate goal. The joint utility (M = 1.72) is just slightly below the integrative level (M = 1.75). The dyad involving both agents with non-specific goals approaches the joint utility (M = 1.11), which is the lowest among the 16 experimental treatments. In addition, when agent goals are homogeneous, highest joint utility (M = 1.65) is achieved when both agents have moderate goals. Furthermore, buying agents achieve the highest individual utility (M = 0.87), when they have a difficult goal while selling agents have a non-specific goal. Similarly, selling agents achieve the highest individual utility (M = 0.88), when they have a difficult goal while buying agents have a non-specific goal. Impasse rate (M = 35.79%) and the number of turns (M = 17.62) are the highest when both agents have difficult goals.

The average performance measures for the treatments with learning are given in Table 9.2. A similar pattern is obtained when compared with Table 9.1. Joint utility is still the highest when selling agents have a difficult goal while buying agents have a moderate goal (M = 1.72). Also, with homogeneous goals, dyad approaches the highest joint utility (M = 1.66) when both agents have moderate goals. In addition, buying agents achieve the highest individual utility (M = 0.87), when they have a difficult goal while selling agents have a non-specific goal. Similarly, selling agents achieve the highest individual utility (M = 0.90), when they have a difficult goal while buying agents have a non-specific goal. Highest impasse rate (M = 36.00%) and highest number of turns (M = 17.81) are observed when both agents have difficult goals.

Table 9.1. Results for Various Combinations of Goal without Learning.

Buying Agent's Goal (b)	Performance Measure	Selling Agent's Goal (s) Non-Specific Mean	Non-Specific Std. Dev.	Easy Mean	Easy Std. Dev.	Moderate Mean	Moderate Std. Dev.	Difficult Mean	Difficult Std. Dev.
Non-Specific	BU(b, s)	0.63	0.05	0.71	0.06	0.73	0.07	0.73	0.08
	SU(b, s)	0.48	0.08	0.76	0.04	0.83	0.02	0.90	0.02
	JU(b, s)	1.11	0.03	1.47	0.03	1.56	0.05	1.63	0.06
	NT(b, s)	6.87	1.21	11.56	1.78	12.91	1.72	16.47	1.32
	IR(b, s)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Easy	BU(b, s)	0.75	0.03	0.78	0.03	0.79	0.04	0.81	0.03
	SU(b, s)	0.56	0.11	0.76	0.04	0.82	0.02	0.88	0.09
	JU(b, s)	1.31	0.08	1.54	0.01	1.61	0.02	1.69	0.02
	NT(b, s)	10.51	1.85	12.72	1.81	13.75	1.70	16.77	1.24
	IR(b, s)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Moderate	BU(b, s)	0.81	0.0156	0.82	0.02	0.83	0.02	0.84	0.02
	SU(b, s)	0.61	0.1107	0.77	0.04	0.82	0.02	0.88	0.01
	JU(b, s)	1.42	0.0967	1.59	0.02	1.65	0.01	1.72	0.01
	NT(b, s)	12.11	1.89	13.61	1.77	14.39	1.66	17.00	1.18
	IR(b, s)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Difficult	BU(b, s)	0.87	0.01	0.86	0.03	0.77	0.13	0.56	0.18
	SU(b, s)	0.72	0.09	0.78	0.05	0.73	0.13	0.56	0.18
	JU(b, s)	1.59	0.09	1.65	0.08	1.50	0.26	1.12	0.35
	NT(b, s)	15.82	1.50	16.40	1.36	16.59	1.32	17.62	1.25
	IR(b, s)	0.00	0.00	0.91	3.83	11.44	14.95	35.79	20.39

Table 9.2. Results for Various Combinations of Goal with Learning.

Buying Agent's Goal (b)	Performance Measure	Selling Agent's Goal (s) Non-Specific Mean	Non-Specific Std. Dev.	Easy Mean	Easy Std. Dev.	Moderate Mean	Moderate Std. Dev.	Difficult Mean	Difficult Std. Dev.
Non-Specific	BU(b, s)	0.69	0.08	0.72	0.06	0.73	0.06	0.73	0.08
	SU(b, s)	0.55	0.19	0.80	0.03	0.84	0.02	0.90	0.02
	JU(b, s)	1.24	0.17	1.52	0.04	1.57	0.04	1.63	0.06
	NT(b, s)	9.79	4.80	13.13	1.83	13.91	1.73	16.69	1.26
	IR(b, s)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Easy	BU(b, s)	0.80	0.02	0.81	0.02	0.81	0.03	0.78	0.02
	SU(b, s)	0.61	0.08	0.77	0.03	0.83	0.02	0.89	0.01
	JU(b, s)	1.41	0.07	1.58	0.02	1.64	0.02	1.67	0.01
	NT(b, s)	12.85	1.95	14.16	1.81	14.75	1.66	19.21	0.32
	IR(b, s)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Moderate	BU(b, s)	0.84	0.01	0.84	0.01	0.84	0.01	0.84	0.01
	SU(b, s)	0.63	0.03	0.79	0.03	0.83	0.02	0.88	0.06
	JU(b, s)	1.62	0.02	1.62	0.02	1.66	0.01	1.72	0.01
	NT(b, s)	14.84	1.72	14.84	1.72	15.24	1.65	17.29	1.12
	IR(b, s)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Difficult	BU(b, s)	0.87	0.01	0.86	0.05	0.77	0.14	0.56	0.18
	SU(b, s)	0.71	0.08	0.78	0.06	0.73	0.13	0.56	0.18
	JU(b, s)	1.58	0.08	1.64	0.10	1.50	0.27	1.11	0.36
	NT(b, s)	16.32	1.38	16.95	1.24	16.93	1.26	17.81	1.22
	IR(b, s)	0.43	1.61	1.49	5.29	12.13	15.56	36.00	20.46

The lowest joint utility (M = 1.11) is, however, obtained when both agents have difficult goals. Learning seems to be effective in improving both individual and joint utilities in the power-effect experiment. Here, a comparison between the learning and no-learning situations indicates that learning is effective on agents with non-specific, easy and moderate goals, but not on agents with difficult goals. Generally, when agents learn, those with more difficult goals achieve higher individual utilities than those with less difficult goals, except in the situation where both agents have difficult goals. Learning seems to help agents to achieve higher individual utilities and higher joint utilities, while the goal difficulty remains in effect in both learning and no-learning environments. With these observations, we now proceed to test the hypotheses given in §§7.4 and 7.5.

9.1 Effect of Goal Setting

Individual Performance

Hypothesis G1. (IR has a Positive and Linear Dependence on Goal Difficulty) We test the effect of goal difficulty on impasse rate by the regression equation:

$$IR = \lambda_0 + \lambda_1 BGOAL + \lambda_2 SGOAL + \varepsilon$$

where IR is the k^{th} replication's average impasse rate, $IR(b, s)^k$, for $b, s \in G$, and BGOAL and SGOAL are b and s, respectively. BGOAL and SGOAL are dummy variables coded as 0 for non-specific, 1 for easy, 2 for moderate, and 3 for difficult. There are 2400 values of IRs (= 150 replications/treatment x 16 treatments), and using these IRs, we find the regression equation as

$$IR = -6.828 + 3.610 BGOAL + 2.947 SGOAL + \varepsilon$$
$$(p < 0.0005) \qquad (p < 0.0005)$$

The value of R^2 of this equation is 0.227, which is low. An explanation might be that impasse rates are often equal to or close to zero, as agents are assigned no goals, easy goals or moderate goals (See Table 9.1). We only observe significant impasse rate (35.79%), in the situation when both agents are assigned difficult goals. Thus, for this setting, IR has a weak linear dependence on goal difficulty on agents with no goals, easy goals or moderate goals. According to the regression equation, the coefficients of BGOAL and SGOAL dummy variables are positive and significant. In addition, the coefficient of the BGOAL variable is greater than that of the SGOAL variable. The buying agent's goal has greater effect on impasse rate than the selling agent's goal. In sum, this hypothesis is supported.

Hypothesis G2. (IU of Specific-Goal Agents > IU of Non-Specific-Goal Agents) Agents with non-specific goals and agents with specific goals are given by:

Agents with Non-Specific Goal:	$\{(b, s) \mid b = \text{Non-Specific}, s \in G\}$	for buyer	(9.1a)
	$\{(b, s) \mid b \in G, s = \text{Non-Specific}\}$	for seller	(9.1b)
Agents with Specific Goal:	$\{(b, s) \mid b \neq \text{Non-Specific}, s \in G\}$	for buyer	(9.2a)
	$\{(b, s) \mid b \in G, s \neq \text{Non-Specific}\}$	for seller	(9.2b)

where each tuple (b, s) in the above equations represents the experimental treatment in Table 9.1, and consists of 150 replications.

To test this hypothesis, we compare the IUs for (9.1a ∪ 9.1b) with those of (9.2a ∪ 9.2b). There are 1200 IUs (= 150 x 4 treatments x 2 for buyer/seller) for non-specific goals and 3600 IUs (= 150 x 12 treatments x 2 for buyer/seller) for specific goals. By conducting a t-test, we observe that agents with specific goals achieve significantly higher individual utilities (M = 0.7912) than those with non-specific goals (M = 0.6456, $p < 0.0005$). Thus, this hypothesis is supported.

Hypothesis G3. (IU has a Positive and Linear Dependence on Goal Difficulty) We compare the individual utilities among agents with easy, moderate and difficult goals. These agents are given by:

Agents with Easy Goals:	$\{(b, s) \mid b = \text{Easy}, s \in G\}$	for buyer	(9.3a)
	$\{(b, s) \mid b \in G, s = \text{Easy}\}$	for seller	(9.3b)
Agents with Moderate Goals:	$\{(b, s) \mid b = \text{Moderate}, s \in G\}$	for buyer	(9.4a)
	$\{(b, s) \mid b \in G, s = \text{Moderate}\}$	for seller	(9.4b)

Agents with Difficult Goals:	$\{(b, s) \mid b = \text{Difficult}, s \in G\}$	for buyer	(9.5a)
	$\{(b, s) \mid b \in G, s = \text{Difficult}\}$	for seller	(9.5b)

There are 600 values of buyer's IUs (= 150 x 4 treatments with each goal), and 600 values of seller's IUs for each of (9.3a $\cup$ 9.3b), (9.4a $\cup$ 9.4b) and (9.5a $\cup$ 9.5b). A one-way ANOVA test indicates significant difference in the individual utilities among these agents, and the results are presented in Table 9.3.

In another case, we modify our test by excluding the treatment in which both agents have difficult goals. This means that the tuple (b, s) for b = Difficult, s = Difficult, is not included in (9.5a $\cup$ 9.5b). The reason to do this is to test if the high impasse rate in the difficult-difficult goal treatment has an effect on this hypothesis. Under this condition, we observe that significant changes do result from this exclusion.

Post-hoc results in Table 9.4 indicate that when the difficult-difficult goal treatment is taken into account, the individual utilities of agents with moderate goals (M = 0.8118) are significantly higher than those of agents with easy goals (M = 0.7764, $p < 0.0005$), but the individual utilities of agents with difficult goals (M = 0.7854) are significantly lower than those of agents with moderate goals (M = 0.7764, $p < 0.0005$). In addition, there is no significant difference in individual utilities between agents with easy and difficult goals. Based on this finding, the hypothesis is only partially supported.

When the difficult-difficult goal treatment is not included, post-hoc results indicate that agents with difficult goals achieve significantly higher individual utilities (M = 0.8608) than agents with moderate goals (M = 0.8118, $p < 0.0005$), and agents with moderate goals also achieve significantly higher individual utilities than those with easy goals (M = 0.7764, $p < 0.0005$).

Table 9.3. Individual Utilities of Agents with Different Goals in No-Learning Environment.

Situation	Agent's Goal	Individual Utility	ANOVA test p-value
All Treatments	Easy Moderate Difficult	0.7764 0.8118 0.7854	< 0.0005
All Treatments except Difficult-Difficult Goal treatment	Easy Moderate Difficult	0.7764 0.8118 0.8608	< 0.0005

Table 9.4. Post-hoc Results for Individual Utilities among Agents with Different Goals in No-Learning Environment.

Situation	Comparison	p-value
All Treatments	Easy vs. Moderate Easy vs. Difficult Moderate vs. Difficult	< 0.0005 0.097 < 0.0005
All Treatments except Difficult-Difficult Goal treatment	Easy vs. Moderate Easy vs. Difficult Moderate vs. Difficult	< 0.0005 < 0.0005 < 0.0005

We further test the impact of agent goals on average individual utility per transaction by the regression equation:

$$AVG\_UTILITY = \lambda_0 + \lambda_1 ROLE + \lambda_2 EGOAL + \lambda_3 MGOAL + \lambda_4 DGOAL + \varepsilon$$

where AVG_UTILITY is the k^{th} replication's average individual utility, $IU(b, s)^k$, for $b, s \in G$, ROLE is the buyer (seller) if IU stands for the buyer's (seller's) IU, and EGOAL, MGOAL and DGOAL are the dummy variables to represent agents with easy, moderate and difficult goal, respectively. Agents with non-specific goal are coded as 0 for all three goal dummy variables. This means that there are 2400 values of buyer's IUs (= 150 x 16 treatments) and 2400 values of seller's IUs. We find the regression equation as

AVG_UTILITY = 0.660 – 0.0279ROLE + 0.131EGOAL + 0.166MGOAL + 0.140DGOAL + ε

($p < 0.0005$) ($p < 0.0005$) ($p < 0.0005$) ($p < 0.0005$)

The equation is significant ($p < 0.0005$, $R^2 = 0.270$), and EGOAL, MGOAL, and DGOAL dummy variables are found to be positive and significant. The value of R^2 is low in this equation, because of the impasse when agents are assigned difficult goals. Since agents receive zero utilities when agents fail to reach agreements, but achieve high utilities if they succeed, IU has a weak linear dependence on goal difficulty, when agents have difficult goals. This is why the coefficient of DGOAL is smaller than that of MGOAL.

Dyadic Performance

Hypothesis G4. (Self and Competitor Goals have Positive and Additive Effects on Joint Utility) We test this hypothesis by the regression equation:

$$\text{JOINT\_UTILITY} = \lambda_0 + \lambda_0\text{BGOAL} + \lambda_0\text{SGOAL} + \lambda_0\text{GOAL\_INT} + \varepsilon$$

where JOINT_UTILITY is the k^{th} replication's average joint utility, $JU(b, s)^k$, for $b, s \in G$, GOAL_INT represents the interaction between the buying agent's goal and the selling agent's goal, and others are defined in the same way as before. There are 2400 values of JUs (= 150 x 16 treatments), and the equation is given by:

JOINT_UTILITY = 0.721 + 0.259BGOAL + 0.303SGOAL – 0.0987GOAL_INT + ε

($p < 0.0005$) ($p < 0.0005$) ($p < 0.0005$)

The equation is significant ($p < 0.0005$, $R^2 = 0.415$), and the coefficients of BGOAL and SGOAL dummy variables are found to be significant and positive. This proposition is supported.

Hypothesis G5. (In Homogeneous Goal Conditions, Highest JU is with Moderate-Moderate Treatment) To test this hypothesis, we compare the joint utilities among the four treatments having homogeneous goal conditions. The results are presented in Table 9.5. We observe that in homogeneous goal conditions, the joint utilities are the highest when both parties have moderate goals. By conducting t-tests, we see that the joint utilities of agreements made at moderate-moderate goal condition are significantly higher (M = 1.6458) than those made at non-specific-non-specific goal condition (M = 1.1115, $p < 0.0005$), easy-easy goal condition (M = 1.5403, $p < 0.0005$), and difficult-difficult goal condition (M = 1.1178, $p < 0.0005$). Overall, this hypothesis is supported.

Table 9.5. Joint Utilities at Different Homogeneous Goal Conditions.

Homogeneous Goal Condition	Joint Utility	t-test p-value (when compared with moderate-moderate goal condition)
No Goal vs. No Goal	1.1115	< 0.0005
Easy vs. Easy	1.5403	< 0.0005
Moderate vs Moderate	1.6458	--
Difficult vs Difficult	1.1178	< 0.0005

Hypothesis G6. (In Disparate Goal Conditions, Highest JU is with Moderate-Difficult Treatment) There are 6 different disparate goal conditions in Table 9.1, and they are:

Non-Specific vs. Easy: {(b, s) | b = Non-Specific, s = Easy}
{(b, s) | b = Easy, s = Non-Specific} (9.6a)

Non-Specific vs. Moderate: {(b, s) | b = Non-Specific, s = Moderate}
{(b, s) | b = Moderate, s = Non-Specific} (9.6b)

Non-Specific vs. Difficult: {(b, s) | b = Non-Specific, s = Difficult}
{(b, s) | b = Difficult, s = Non-Specific} (9.6c)

Easy vs. Moderate: {(b, s) | b = Easy, s = Moderate}
{(b, s) | b = Moderate, s = Easy} (9.6d)

Easy vs. Difficult: {(b, s) | b = Easy, s = Difficult}
{(b, s) | b = Difficult, s = Easy} (9.6e)

Moderate vs. Difficult: {(b, s) | b = Moderate, s = Difficult}
{(b, s) | b = Difficult, s = Moderate} (9.6f)

There are 300 values of JUs (= 150 x 2 treatments) for each of (9.6a) through (9.6f), and the average joint utilities for these conditions are presented in Table 9.6.

The highest joint utilities (M = 1.7151) are achieved when the selling agents have a difficult goal while the buying agents have a moderate goal. However, the joint utilities are much lower (M = 1.5053) when the selling agents have a moderate goal while the buying agents have a difficult goal. We further observe that the joint utilities of the moderate-difficult goal group are only significantly higher (M = 1.6102) than those of the non-specific-easy goal group (M = 1.3888, $p < 0.0005$) and those of the non-specific-moderate goal group (M = 1.4904, $p < 0.0005$). However, the joint utilities of the moderate-difficult goal group are significantly lower than those of the easy-difficult goal group (M = 1.6703, $p < 0.0005$). Furthermore, insignificant difference in joint utilities is found between moderate-difficult goal group and non-specific-difficult goal group, as well as between moderate-difficult goal and easy-moderate goal group. Thus, this hypothesis is only partially supported.

Table 9.6. Joint Utilities at Different Disparate Goal Conditions.

Disparate Goal Condition	**Joint Utility**	**t-test p-value** (vs. moderate-difficult goal condition)	**t-test p-value** (vs. easy-difficult goal condition)
No Goal vs. Easy	1.3888	< 0.0005	< 0.0005
No Goal vs. Moderate	1.4904	< 0.0005	< 0.0005
No Goal vs. Difficult	1.6085	1.000	< 0.0005
Easy vs. Moderate	1.6006	0.901	< 0.0005
Easy vs. Difficult	1.6703	< 0.0005	--
Moderate vs. Difficult	1.6102	--	< 0.0005

On the other hand, when the easy-difficult goal condition is compared with other disparate goal conditions by t-tests, results from Table 9.6 indicate that the joint utilities of the easy-difficult goal group are significantly higher than those of the other five disparate goal groups ($p < 0.0005$ for all comparisons).

Hypothesis G7. (Self and Competitor Goals have Positive and Additive Effects on NT) We test the impact of agent goals on the average number of turns by the regression equation:

$$AVG\_TURNS = \lambda_0 + \lambda_1 BGOAL + \lambda_2 SGOAL + \lambda_3 GOAL\_INT + \varepsilon$$

where AVG_TURNS is the k^{th} replication's average number of turns, $NT(b, s)^k$, for $b, s \in G$, and the other variables are defined in the same way as before. This means that there are 2400 values of NTs (= 150 x 16 treatments). The equation is:

$$AVG\_TURNS = 1.056 + 3.430BGOAL + 3.724SGOAL - 0.780GOAL\_INT + \varepsilon$$

$(p < 0.0005)$ $(p < 0.0005)$ $(p < 0.0005)$

The equation is significant ($p < 0.0005$, $R^2 = 0.736$), and the coefficients of both BGOAL and SGOAL dummy variables are significant and positive. Thus, this proposition is supported.

Relative Effects of Goals

Hypothesis G8. (In Disparate Goal Conditions, IU of More-Difficult-Goal Agents > IU of Less-Difficult-Goal Agents) We conducted a chi-square test to investigate the significance of the difference of the individual utilities between agents with disparate goals. There are a total of 12 disparate goal treatments, and a total of 1800 counts in our experiment. Among 1800 cases, there are 1609 cases that support the proposition. The critical value for the χ^2 distribution with one degree of freedom and 95% confidence level is 3.84. The value of χ^2 in our experimental setting is found to be 1117.07, which is greater than the critical value. As a result, this hypothesis is supported.

9.2 Joint Effect of Goal Setting and Learning

Individual Performance

Hypothesis GL1. (IU of Agents with Learning > IU of Agents without Learning) To test this proposition, we compare the individual utilities of learning agents with those of non-learning ones. These agents are given by:

$$\text{Agents without Learning:} \quad \{(b, s) \mid b, s \in G\} \tag{9.7}$$

$$\text{Agents with Learning:} \quad \{(b, s) \mid b, s \in GL\} \tag{9.8}$$

There are 4800 values of IUs (= 150 x 32 treatments – 16 without learning and 16 with learning) for each of (9.7) and (9.8), in which 2400 values are buyer's IUs and the other 2400 values are seller's IUs.

We conducted a t-test to compare the individual utilities between agents with and without learning. Results indicate that agents with learning achieve significantly higher individual utilities ($M = 0.7680$) than those without learning ($M = 0.7548$, $p < 0.0005$). We further test the impact of agent goals and learning on average individual utility per transaction by the regression equation:

$$\text{AVG\_UTILITY} = \lambda_0 + \lambda_1\text{ROLE} + \lambda_2\text{GOAL} + \lambda_3\text{LEARNING} + \lambda_4\text{GOAL\_LEARNING\_INT} + \varepsilon$$

where AVG_UTILITY is the k^{th} replication's average individual utility, $\text{IU}(b, s)^k$, for $b, s \in G \cup GL$, GOAL represents $b(s)$ for buyers (sellers), LEARNING indicates whether an agent learns or not, GOAL_LEARNING_INT is the interaction between goal and learning, and other variables are defined in the same way as discussed above. We find the equation as:

$$\begin{aligned}\text{AVG\_UTILITY} = 0.621 &- \underset{(p < 0.0005)}{0.0271}\text{ROLE} + \underset{(p < 0.0005)}{0.05368}\text{GOAL} + \underset{(p < 0.0005)}{0.03377}\text{LEARNING} + \\ &+ \underset{(p < 0.0005)}{0.00821}\text{GOAL\_LEARNING\_INT} + \varepsilon\end{aligned}$$

The LEARNING variable is found to be significant and positive. However, we achieve a low value of R^2. These are two reasons for this. First, there is a high impasse rate when both the buying and selling agents are assigned difficult goals. Since agents receive zero utilities when agents fail to reach agreements, but achieve high utilities if they succeed, IU has a weak linear dependence on goal difficulty, when agents have difficult goals. Second, agents have different concession behaviors in our experiments. The variation in β has an impact on the agent's negotiation performance. In sum, this hypothesis is supported.

Hypothesis GL2. (Less-Difficult-Goal Agents have Higher Increase in Utility from Learning) First, we investigate if learning can improve individual utilities for agents with different goals. We compare the individual utilities between agents with and without learning, when agents have different assigned goals. Agents with learning and different goals are given by:

Non-Specific-Goal Agents with Learning:	$\{(b, s) \mid b =$ Non-Specific with Learning, $s \in GL\}$	for buyer	(9.9a)
	$\{(b, s) \mid b \in GL, s =$ Non-Specific with Learning$\}$	for seller	(9.9b)
Easy-Goal Agents with Learning:	$\{(b, s) \mid b =$ Easy with Learning, $s \in GL\}$	for buyer	(9.10a)
	$\{(b, s) \mid b \in GL, s =$ Easy with Learning$\}$	for seller	(9.10b)
Moderate-Goal Agents with Learning:	$\{(b, s) \mid b =$ Moderate with Learning, $s \in GL\}$	for buyer	(9.11a)
	$\{(b, s) \mid b \in GL, s =$ Moderate with Learning$\}$	for seller	(9.11b)
Difficult-Goal Agents with Learning:	$\{(b, s) \mid b =$ Difficult with Learning, $s \in GL\}$	for buyer	(9.12a)
	$\{(b, s) \mid b \in GL, s =$ Difficult with Learning$\}$	for seller	(9.12b)

where each tuple (b, s) in the above equation represents the experimental treatment in Table 9.2.

We conducted t-tests to compare the individual utilities between (9.1a $\cup$ 9.1b) and (9.9a $\cup$ 9.9b), (9.3a $\cup$ 9.3b) and (9.10a $\cup$ 9.10b), (9.4a $\cup$ 9.4b) and (9.11a $\cup$ 9.11b), and (9.5a $\cup$ 9.5b) and (9.12a $\cup$ 9.12b). The results are presented in Table 9.7. We observe that agents with non-specific goals, easy goals, and moderate goals achieve significantly higher individual utilities when they learn. However, for agents with difficult goals, we do not find any significant difference in individual utilities between those without learning and those with learning.

Table 9.7. Effect of Learning on Individual Utilities at Different Goals.

Goal	Individual Utility		Average Increase	t-test p-value
	Without Learning	With Learning		
No Goal	0.6456	0.6709	0.0254	< 0.0005
Easy	0.7764	0.7928	0.0164	< 0.0005
Moderate	0.8118	0.8235	0.0117	< 0.0005
Difficult	0.7854	0.7849	-0.0004	0.950

Next, we test if the *increases* in utilities are different among the goal conditions. We conducted a 2 x 4 two-way ANOVA test that consists of agents with and without learning at four different goal conditions. We compare the *increases* in individual utilities between (9.1a ∪ 9.1b) and (9.9a ∪ 9.9b), (9.3a ∪ 9.3b) and (9.10a ∪ 9.10b), (9.4a ∪ 9.4b) and (9.11a ∪ 9.11b), and (9.5a ∪ 9.5b) and (9.12a ∪ 9.12b). The results are presented in Figure 9.1, which shows an interaction effect between goal difficulty and learning.

According to Figure 9.1, when agents learn, agents with non-specific goals achieve higher increases in individual utility (M = 0.0254) than those with easy goals, moderate goals and difficult goals (M = 0.0164, M = 0.0117, M = -0.0004 respectively). We observe that learning provides a higher increase in individual utility for agents with less difficult goals, as compared to agents with more difficult goals. Thus, this hypothesis is supported.

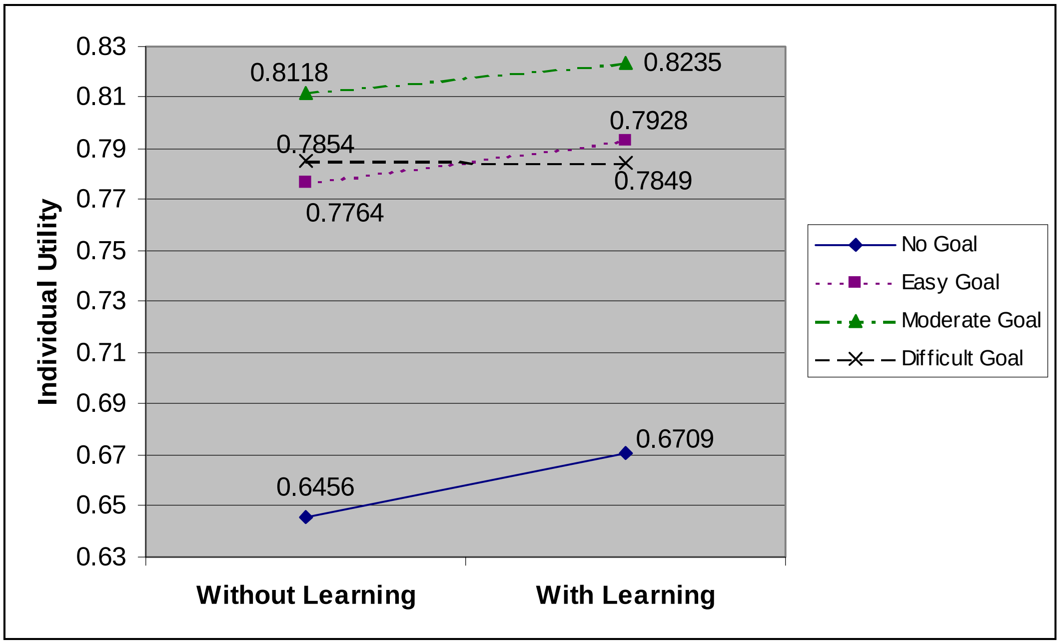

Figure 9.1. Interaction Effect of Goal Setting and Learning on Individual Utility.

Hypothesis GL3. (When Agents Learn, IU of More-Difficult-Goal Agents > IU of Less-Difficult-Goal Agents) To test this hypothesis, we compare the individual utilities among agents with easy goals and learning (9.10a ∪ 9.10b), agents with moderate goals and learning (9.11a ∪ 9.11b), and agents with difficult goals and learning (9.12a ∪ 9.12b). This comparison is also done by including the difficult-difficult treatment from (9.12a ∪ 9.12b) in one case, and excluding it in another case. The results are presented in Tables 9.8 and 9.9.

Table 9.8. Individual Utilities of Agents with Different Goals in Learning Environment.

Situation	Agent's Goal	Individual Utility	One-way ANOVA test p-value
All Treatments	Easy Moderate Difficult	0.7928 0.8235 0.7849	< 0.0005
All Treatments except Difficult-Difficult Goal treatment	Easy Moderate Difficult	0.7928 0.8235 0.8608	< 0.0005

Table 9.9. Post-hoc Results for Individual Utilities of Agents with Different Goals in Learning Environment.

Situation	Goal Comparison	p-value
All Treatments	Easy vs. Moderate Easy vs. Difficult Moderate vs. Difficult	< 0.0005 0.281 < 0.0005
All Treatments except Difficult-Difficult Goal treatment	Easy vs. Moderate Easy vs. Difficult Moderate vs. Difficult	< 0.0005 < 0.0005 < 0.0005

When the treatment consisting of difficult-difficult goal condition is taken into account, results indicate that when agents learn, the individual utilities of agents with moderate goals (M = 0.8235) are significantly higher than those of agents with easy goals (M = 0.7928, $p < 0.0005$), but the individual utilities of agents with difficult goals (M = 0.7849) are significantly lower than those of agents with moderate goals (M = 0.7928, $p < 0.0005$). In addition, there is no significant difference in individual utilities between agents with easy and difficult goals. Based on this finding, this hypothesis is only partially supported.

However, when the difficult-difficult treatment is taken out, this hypothesis is supported. Results from Tables 9.8 and 9.9 indicate that when agents learn, agents with difficult goals achieve significantly higher individual utilities (M = 0.8608) than agents with moderate goals (M = 0.8235, $p < 0.0005$), while agents with moderate goals achieve significantly higher individual utilities than agents with easy goals (M = 0.7928, $p < 0.0005$).

Dyadic Performance

Hypothesis GL4. (JU of Agreements with Learning > JU of Agreements without Learning) To test this proposition, we conducted a t-test to compare the joint utilities of 16 experimental treatments in Table 9.1 with those of 16 experimental treatments in Table 9.2. They are given by the $JU(b, s)^k$ for $b, s \in G$, and the $JU(b, s)^k$ for $b, s \in GL$. Results indicate that the joint utilities of dyads with learning are significantly higher (M = 1.5361) than those of dyads without learning (M = 1.5096, $p < 0.0005$). Thus, this hypothesis is supported.

9.3 Discussion

9.3.1 Effect of Goal Setting

A summary of the results of the effect of goal setting on electronic negotiation outcome is presented in Table 9.10, wherein the shaded portion indicates hypotheses that are not fully supported. Generally, agents with specific goals perform better than those with non-specific goals. Moreover, agents with more difficult goals achieve higher individual utilities than those with less difficult goals. However, lower performance is achieved when both parties have difficult goals, because of a high impasse rate. Since the impasse rate has a positive and linear dependence on goal difficulty, we recommend agents to be assigned goals, but they should not be assigned difficult goals in all cases.

To approach highly integrative agreements, we need to consider the goals of both parties. In homogeneous goal conditions, joint utilities are the highest when both agents have moderate goals. Difficult-difficult goal combination is not recommended. In heterogeneous goal conditions, our results are only partially consistent with the findings from Huber & Neale (1986), because of the differences in payoff tables between the buying and the selling agents. To compute the utilities, buying agents incorporate the values of both the negotiating and the non-negotiating issues, while selling agents consider the negotiating issues only. There is only one type of product for buying agents to choose, while selling agents can sell their products to any buying agent. Thus, a same value of *final threshold* for both a buying and a selling agent means that achieving certain utility level is more difficult for the buying agent than for the selling agent. This is also the reason why the goals of buying agents have a greater impact than the goals of selling agents on the impasse rates. As a result, when buying agents have difficult goals, assigning easy goals to selling agents can maximize the joint utilities of the agreements. On the other hand, when selling agents have difficult goals, we can make the agreements to be more integrative by assigning moderate goals to buying agents.

Although agents with more difficult goals achieve higher individual utilities than those with less difficult goals, longer time is required to reach agreements. Thus, the trade-off between performance and time is a considerable factor, when we assign goals to negotiation agents.

Table 9.10. Summary of the Effects of Goal Setting on Electronic Negotiation Outcome.

	Hypothesis	Our Results	Literature supporting Hypothesis	Literature not supporting Hypothesis
G1	Impasse rate has a positive and linear dependence on goal difficulty.	Supported	Pruitt (1981)	None
G2	Agents with specific goals will have higher individual utilities than those with non-specific goals.	Supported	Huber & Neale (1986)	None
G3	Individual utility has a positive and linear dependence on goal difficulty.	Partially Supported	Hamner & Harnett (1974) Bazerman & Neale (1985) Huber & Neale (1987) Huber & Neale (1986) partially	None
G4	The assigned goals of both agents will have a positive and additive effect on joint utility.	Supported	Huber & Neale (1986)	None
G5	When agent goals are homogeneous, joint utility will be the highest when both agents have moderate goals.	Supported	Huber & Neale (1986)	None
G6	When agent goals are disparate, joint utility will be the highest when one agent has a difficult goal and the other has a moderate goal.	Partially Supported	Huber & Neale (1986)	None
G7	The assigned goals of both agents will have a positive and additive effect on the time to reach agreements.	Supported	Pruitt (1981)	None
G8	When goals are disparate between two agents in a negotiation, the agent with the more difficult goals will achieve higher individual utility than the one with the less difficult goals.	Supported	Huber & Neale (1986)	None

9.3.2 Joint Effect of Goal Setting and Learning

Table 9.11 presents the summary of the results of the joint effect of goal setting and learning on electronic negotiation outcome. Our heuristic algorithms for learning are found to be effective for agents with non-specific, easy, and moderate goals. However, agents with difficult goals do not achieve higher individual utilities when they learn. In addition, when agents learn, agreements become more integrative. In general, learning is beneficial to agents with no goals, easy goals or moderate goals, but it is not effective for agents with difficult goals.

We also observe that learning provides a higher increase in individual utility for agents with less difficult goals, when compared to agents with more difficult goals. This implies that learning is a more important factor for agents with non-specific and easy goals than for agents with moderate and difficult goals for consideration in electronic negotiation.

When learning is implemented and both agents are not assigned difficult goals, we observe that agents with more difficult goals achieve higher individual utilities than those with less difficult goals. That is, although we find that learning is effective in improving the individual performance of an agent, goal difficulty remains in effect when agents learn. Thus, assigning goals to agents is useful, even when agents learn.

Table 9.11. Summary of the Joint Effect of Goal Setting and Learning on Electronic Negotiation Outcome.

	Hypothesis	Our Results	Literature supporting Hypothesis	Literature not supporting Hypothesis
GL1	Agents with learning will have higher individual utilities than those without learning.	Supported	Zeng & Sycara (1998) Deveaux et al. (2001)	None
GL2	Learning will provide a higher increase in individual utility for agents with less difficult goals, as compared to agents with more difficult goals.	Supported	None	None
GL3	When agents learn, the agents with more difficult goals will achieve higher individual utilities than those with less difficult goals.	Partially Supported	None	None
GL4	Joint utility with learning will be higher than that in negotiation without learning.	Supported	None	None

10 Conclusion

The broad objective of this thesis is to develop methods for negotiation agents to understand their owners' preferences and to learn their opponents' behavior, so that a human being's interaction is minimized and negotiation performance can be improved over time. Research on incorporating human factors into negotiation agents is at a preliminary stage. We attempt to make negotiation agents to be more similar to human negotiators, by integrating learning and two psychological factors – situational power and goal constraints – into negotiation agents.

Learning an opponent's negotiation behavior is important, because the expected concession can influence a negotiator's behavior (Pruitt, 1981; Raiffa, 1982). The learning algorithms we have presented are based on the premise that negotiation experience of opponents is missing. This is in contrast to genetic-algorithm-based learning and Bayesian learning, in both of which agents have some prior knowledge and beliefs of their opponents. With incomplete information about an opponent's concession behavior, we can use an opponent's initial demand level and concession rate as only a guidance. In this thesis, we have presented heuristic algorithms to estimate the parameters of time-dependent tactic, and then to react to the estimates, so as to achieve higher performance. Experimental results indicate that except for agents with difficult goals, our heuristic algorithms for learning are effective in improving both individual and dyadic performances.

Besides learning, we have presented how AHP can be used to incorporate situational power into negotiation agents. This idea is similar to AHP's applicability in benchmarking an organization's performance (Forman & Gass, 2001), but it is new to the negotiation context. Experimental results indicate that the outcomes in electronic negotiation agree with those observed in human negotiation. Agents with higher power achieve significantly higher performance than those with lower power, but take longer time to reach agreements. The trade-off between performance and time is a considerable factor, when we integrate situational power into electronic negotiation. Moreover, we find that our learning algorithms are more effective for high-power agents than for low-power agents in electronic negotiation.

In addition, we have shown how to incorporate goal constraints into negotiation agents by asking users to specify threshold values that agents must achieve before concluding the negotiations. With the effect of goal constraints on human negotiation as a reference, we achieve consistent results in electronic negotiation. Agents with specific and difficult goals achieve higher performance than those with easy or no goals, but attain higher impasse rates, and require longer time for reaching agreements. Thus, negotiators should not assign difficult goals in all cases, because there is a higher chance of failing to reach agreements. Also, as with situational power, the trade-off between performance and time must be considered. Furthermore, for the joint effect of learning and goal constraints, we observe that our heuristic algorithms for learning are effective for agents with non-specific, easy, and moderate goals. However, agents with difficult goals do not achieve higher performance when they learn. Learning is a more important factor for agents with less difficult goals than those with more difficult goals.

In summary, our work provides three major contributions to the field of information systems and electronic-commerce research, in which electronic negotiations are frequently studied. First, we have designed algorithms for learning the time-dependent tactic of an opponent. Second, we have determined that our learning algorithms are beneficial to negotiation agents. Our third contribution, which is relative minor, is to present how learning, situational power and goal constraints can be integrated into an agent. Finally, with the integration of these factors into an agent, we have observed that there is a congruence of outcome between human and electronic negotiations. This congruence indicates that our integration methodologies are worthwhile to consider.

Several extensions to our study are possible. First, since one of the objectives of this thesis is to assess the level of concurrence between the outcomes from human and electronic negotiations, all the settings in our experiments are based on earlier studies on human negotiation. In our experiments, all agents were only allowed to negotiate with a fixed T_{max}. However, different users or agents are more likely to have different negotiation time limits in the real world. Besides this, other parameters such as P_{min}, P_{max} and K are going to be varied for different people. To obtain more meaningful results that represent the real-world situation, it is important to conduct further research and experiments with agents having varying values of the above-mentioned parameters. This could further validate our findings in other situations.

In addition, our experiments focus on the negotiation of one product whose properties are undifferentiable. In an electronic-commerce business, products from different agents are usually similar instead of being exactly the same as one another. Difference of preferences on similar-quality products might affect negotiation outcome. Thus, further experiments involving more than one product, but with similar properties, is important to undertake.

Furthermore, McAlister *et al.* (1986) conducted human negotiation experiment to identify the joint effect of situational power and goal constraints. Their results indicated that for negotiators with low power, having moderately difficult goals led to lower performance than having no goals. On the other hand, for negotiators with equal or high power, having moderately difficult goals led to higher performance than having no goals. One aim for the future is to validate this joint effect in electronic negotiation.

Finally, in this thesis, we focus on investigating how situational power and goal constraints can affect the negotiation behavior of an agent itself. Besides understanding an owner's preferences, detecting an opponent's situational power and goal constraints might be beneficial. This requires further research.

Bibliography

Axelrod R. (1984). *The Evolution of Cooperation*. Basic Books, Inc., Publishers, New York.

Bazerman, M. H., Magliozzi, T., & Neale, M. A. (1985). Integrative Bargaining in a Competitive Market. *Organizational Behavior and Human Decision Processes, 35*, 194-313.

Chatterjee, S. & Price, B. (1991). *Regression Analysis by Example 2nd Ed*. John Wiley & Sons, Inc., New York.

Choi, S. P. M., Liu, J., & Chan, S. P. (2001). A Genetic Agent-Based Negotiation System. *Computer Networks, 37*, 195-204.

Deveaux, L., Paraschiv, C., & Latourrette, M. (2001). Bargaining on an Internet Agent-Based Market: Behavioral vs. Optimizing agents. *Electronic Commerce Research, 1*, 371-401.

Dwyer, R. F., & Walker, O. C. Jr. (1981). Bargaining in an Asymmetrical Power Structure. *Journal of Marketing, 45 (Winter)*, 104-115.

Eliashberg, J., LaTour, S. A., Rangaswamy, A., & Stern, L. W. (1986). Assessing the Predictive Accuracy of Two Utility-Based Theories in a Marketing Channel Negotiation Context. *Journal of Marketing Research, 23 (May)*, 101-110.

Emerson, R. M. (1962). Power-Dependence Relations. *American Sociological Review, 27 (February)*, 31-40.

Faratin P., Sierra, C., & Jennings N. R. (1997). Negotiation Decision Functions for Autonomous Agents. *Robotics and Autonomous Systems, 24 (3-4)*, 159-182.

Forman, E. H., & Gass, S. I. (2001). The Analytical Hierarchy Process – An Exposition. *Operations Research, 49 (4)*, 469-486.

Ghodsypour, S. H., & O'Brien, C. (1998). A Decision Support System for Supplier Selection using an Integrated Analytic Hierarchy Process and Linear Programming. *Int. J. Production Economics, 56-57*, 199-212.

Goh, K. Y., Teo, H. H., Wu, H., & Wei, K. K. (2000) Computer-Supported Negotiations: An Experimental Study of Bargaining in Electronic Commerce. Proceedings of the Twenty First International Conference on Information Systems.

Greenhalgh, L., Neslin, S. A., & Gilkey, R. W. (1985). The Effects of Negotiator Preferences, Situational Power, and Negotiator Personality on Outcomes of Business Negotiations. *Academy of Management Journal, 28 (1)*, 9-33.

Hamner, W.C., & Harnett, D. L. (1974). Goal Setting, Performance and Satisfaction in an Interdependent Task. *Organization Behavior and Human Performance*, *12*, 217-230.

Huber, V. L., & Neale, M. A. (1986). Effects of Cognitive Heuristics and Goals on Negotiator Performance and Subsequent Goal Setting. *Organizational Behavior and Human decision Processes, 36*, 342-365

Huber, V. L., & Neale, M. A. (1987). Effects of Self- and Competitor Goals on Performance in an Interdependent Bargaining Task. *Journal of Applied Psychology, 72 (2)*, 197-203.

Impulse: MIT Media Lab (1999). http://agents.www.media.mit.edu/groups/agents/projects/impulse/

Jennings, N. (2001). Agent-Mediated Electronic Commerce http://www.ecs.soton.ac.uk/~nrj/amec.html

Kasbah: MIT Media Lab (1999). http://ecommerce.media.mit.edu/Kasbah/

Kofler, M. (1997). Maple: *An Introduction and Reference*. Addison Wesley Longman

Lim, L., & Benbasat, I. (1992). A Theoretical Perspective of Negotiation Support Systems. *Journal of Management Information Systems, 9 (3)*, 27-44

Locke, E. A., Shaw, K. N., Saari, L. M., & Latham, G. P. (1981). Goal Setting and Task Performance 1969-1980. *Psychological Bulletin, 90*, 125-152.

Luce, R. D., & Raiffa, H. (1957). *Games and Decisions*; *Introduction and Critical Survey*. New York: Wiley.

Matos, N., Sierra, C., & Jennings, N. R. (1998). Determining Successful Negotiation Strategies: An Evolutionary Approach. *Proc. 3rd Int. Conf. on Multi-Agent Systems (ICMAS-98)*, Paris, France, 182-189.

McAlister, L., Bazerman, M. H., & Fader, P. (1986). Power and Goal Setting in Channel Negotiations. *Journal of Marketing Research, 23*, 228-237.

Murugesan, S. (1998). Intelligent Agents on the Internet and Web, *TENCON '98. 1998 IEEE Region 10 International Conference on Global Connectivity in Energy, Computer, Communication and Control, 1 (December),* 97-102.

Murugesan S. (2000). Negotiation by Software Agents in Electronic Marketplace. *IEEE*, 286-290.

Nash, J. (1950). The Bargaining Problem. *Econometrica, 18*, 155-162.

Nash. J. (1953). Two-Person Cooperative Games. *Econometrica, 21*, 129-140.

Neslin, S. A., & Greenhalgh L. (1983). Nash's Theory of Cooperative Games as a Predictor of the Outcomes of Buyer-Seller Negotiations: An Experiment in Media Purchasing. *Journal of Marketing Research, 20 (November)*, 368-379.

Park, S., Durfee, E. H., & Birmingham, W. P. (1999). An Adaptive Agent Bidding Strategy based on Stochastic Modelling. *ACM*, 147-153.

Pinkley, R. L., Neale, M. A., & Bennett, R. J. (1994). The Impact of Alternatives to Settlement in Dyadic Negotiation. Organizational Behavior and Human Decision Processes, 57, 97-116.

Pruitt, D. G. (1981). *Negotiation Behavior.* New York: Academic Press.

Raiffa, H. (1982). *The Art and Science of Negotiation*. Boston: Harvard University Press.

Ratkowsky D. A. (1983). *Nonlinear Regression Modeling: A Unified Practical Approach.* Marcel Dekker, Inc., New York.

Rubin, J. Z., & Brown, B. R. (1975). *The Social Psychology of Bargaining and Negotiation.* Academic Press.

Sim, K. M., & Chan, R. (2000). A brokering protocol for agent-based e-commerce. *IEEE Transactions on Systems, Man, and Cybernetics – Part C: Applications and Reviews, 30 (4)*, 474-484.

Summer, G. W., Peters, W. S., & Armstrong, C. (1977). *Basic Statistics in Business and Economics, 2nd Ed*. Wadsworth Publishing Company, Inc. Belmont.

Sundarraj, R. (2002). *Preference Elicitation and its Effects on Electronic Shopping Agents*. A Working Paper at the Department of Management Science, University of Waterloo.

Teich, J., Wallenius, H., & Wallenius, J. (1999). Multiple-issue Auction and Market Algorithms for the World Wide Web. *Decision Support Systems, 26*, 49-66.

Tripp, T. M., & Sondak, H. (1992). An Evaluation of Dependent Variables in Experimental Negotiation Studies: Impasse Rates and Pareto Efficiency. *Organizational Behavior and Human Decision Processes, 51*, 273-295.

Tu, M. T., Wolff, E., & Lamersdorf, W. (2000). Genetic Algorithms for Automated Negotiations: A FSM-Based Application Approach. *IEEE*, 1029-1033.

Ye, Y., Liu, J., & Moukas, A. (2001). Agents in Electronic Commerce. *Electronic Commerce Research, 1*, 9-14.

Zacharia, G., Evgeniou, T., Moukas, A., Boufounos, P., & Maes, P. (2001). Economics of Dynamic Pricing in a Reputation Brokered Agent Mediated Marketplace. *Electronic Commerce Research*, *1*, 85-100.

Zeng, D., & Sycara, K. (1998). Bayesian learning in negotiation. *International Journal of Human-Computer Studies*, *48*, 125-141.