Reputation based Buyer Strategies for Seller Selection in Electronic Markets

BY

C2008
Sandhya Beldona

Submitted to the Department of Electrical Engineering & Computer Science and the Graduate Faculty of the University of Kansas in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

____
Dr. Arvin Agah

____
Dr. Costas Tsatsoulis

____
Dr. Gunes Ercal

____
Dr. Luke Huan

____
Dr. Prasad Anil Kulkarni

____
Dr. Prakash P Shenoy

Date Defended________________________________

The Dissertation Committee for Sandhya Beldona certifies
that this is the approved version of the following dissertation:

Reputation based Buyer Strategies for Seller Selection in Electronic Markets

Committee:

____
Dr. Arvin Agah

____
Dr. Costas Tsatsoulis

____
Dr. Gunes Ercal

____
Dr. Luke Huan

____
Dr. Prasad Anil Kulkarni

____
Dr. Prakash P Shenoy

Date approved: ______________________________

Acknowledgments

This dissertation would not have been possible without the support of faculty, staff, and friends at University of Kansas, and the support and love of my family and friends.

First, I would like to thank my advisor, Dr. Tsatsoulis for his time, patience, support, and encouragement. I had been very apprehensive about starting my Ph.D. as it meant staying in a different town, away from my husband and children for an extended period of time. I also needed financial support to cover my living expenses during my studies. At this juncture, I met Dr. Tsatsoulis. He was very encouraging and understanding of my situation. I was very fortunate that he offered me research assistantship for doing research in the area of intelligent systems and multi agents systems which was my preferred interest. Since then he has been my mentor. He has patiently guided me in developing my research abilities, critical thinking, and research writing skills. His belief in me made me continue when my confidence was wavering. I have learnt a lot from him and will be always be grateful for that.

I would like to thank Dr. Agah for agreeing to serve as co chair of my dissertation committee. He has always been very kind and encouraging. His comments and suggestions have helped in improving the quality of my research and dissertation. I will always remember with gratitude, his help on the day I defended.

I would like to thank my committee members: Dr. Ercal, Dr. Gill, Dr. Huan, Dr. Kulkarni, and Dr. Shenoy. Their comments helped in improving the quality of my research and dissertation. They also gave me numerous suggestions to improve quality of my job talk for academic positions. I am grateful to Dr. Ercal, Dr. Gill, Dr. Huan, and Dr Kulkarni for agreeing to be on my committee in the final stages of my dissertation. Dr. Ercal and Dr. Shenoy have been extremely helpful.

I would like to thank the department staff, especially Paula Conlin, and Pamela Shadoin for their help in making my graduate student life progress smoothly.

Finally and most importantly, I would like to thank my family members and friends. This would not have been possible for me without the continuous support and love of my husband, Sriram. He motivated me to embark on my Ph.D. journey. He gave me moral support when the journey seemed difficult. He took care of our children enabling me to study without interruptions. My parents offered me moral support from India. My children, Varun and Varsha, have been very patient and understanding throughout. In spite of being very young, they adjusted very well and remained cheerful when living in a different town without me. They always cheered me to do well in my exams. I would also like to thank all my friends in Wichita. They helped me by encouraging me to continue, and took care of my children when needed.

Abstract

Reputation based adaptive buying agents that reason about sellers for purchase decisions have been designed for B2C ecommerce markets. Previous research in the area of buyer agent strategies for choosing seller agents in ecommerce markets has focused on frequent purchases. In this thesis, we present reputation based strategies for buyer agents to choose seller agents in a decentralized multi agent based ecommerce markets for frequent as well as infrequent purchases.

We consider a marketplace where the behavior of seller agents and buyer agents can vary, they can enter and leave the market any time, they may be dishonest, and quality of the product can be gauged after actually receiving the product. Buyer agents exchange seller agents' information, which is based on their own experiences, with other buyer agents in the market. However, there is no guarantee that when other buyer agents provide information, they are truthful or share similar opinions.

First we present a method for buyer agent to model a seller agent's reputation. The buyer agent computes a seller agent's reputation based on its ability to meet its expectations of product quality and price as compared to its competitors. We show that a buying agent acting alone, utilizing our model of maintaining seller agents' reputation and buying strategy does better than buying agents acting alone employing strategies proposed previously by other researchers for frequent as well as for infrequent purchases.

Next we present two methods for buyer agents to identify other trustworthy buyer agent friends who are honest and have similar opinions regarding seller agents, based on sharing of seller agents' information with each other. In the first method, buyer agent utilizes other buyer agents' opinions and ratings of seller agents to identify trustworthy buyer agent friends. Reputation of seller agents provided by trustworthy buyer agent friends is adjusted to account for the differences in the rating systems and combined with its own information on seller agents to choose high quality, low priced seller agent. In the second method, buyer agent only utilizes other buyer agents' opinions of seller agents to identify trustworthy buyer agent friends. Ratings are assigned to seller agents by the buyer agent based on trustworthy friend buyer agents' opinions and combined with its own rating on seller agents to choose a high quality, low priced seller agent to purchase from.

We conducted experiments to show that both methods are successful in distinguishing between trustworthy buyer agent friends, whose opinions should be utilized in decision making, and untrustworthy buyer agent friends who are either dishonest, or have different opinions. We also show that buyer agents using our models of identifying trustworthy buyer agent friends have higher performance than a buyer agent acting alone for infrequent purchases and for increasing numbers of sellers in the market.

Finally we analyze the performances of buyer agents with risk taking and conservative attitudes. A buyer agent with risk taking attitude considers a new seller agent as reputable initially and tends to purchase from a new seller agent if they are offering the lowest price among reputable seller agents. A buyer agent with conservative attitude is cautious in its approach and explores new seller agents at a rate proportional to the ratio of unexplored seller agents to the all the seller agents who have sent bids. Our results show that, when buyer agents are making decisions based on their own information, a buyer agent with conservative attitude has the best performance. When buyer agents are utilizing information provided by their trusted friends, a buyer agent with risk taking attitude and using only trusted friend buyer agents' opinions of seller agents has the best performance.

In summary, the main contributions of this dissertation are:

1. A new reputation based way to model seller agents by buyer agents based on direct interactions.

2. A protocol to exchange reputation information about seller agents with other buyer agent friends based on the friends' direct interaction with seller agents.

3. Two methods of identifying trustworthy buyer agent friends who are honest and share similar opinions, and utilizing the information provided by them to maximize a buyer agent's chances of choosing a high quality, low priced seller agent to purchase from.

Contents

List of Figures

List of Tables

Chapter 1 Introduction

1.1 Motivation

The Electronic Commerce Association defines *ecommerce* as "*electronic commerce [that] covers any form of business or administrative transaction or information exchange that is executed using any information and communications technology*" [48]. It is a very broad definition covering activities like purchasing goods over the telephone or using credit cards to make purchases. In this dissertation, *ecommerce* is used to refer to the buying and selling of goods on the Internet. According to the Consumer Buying Behavior (CBB) model which covers Business to Consumer (B2C) transactions [33], there are six fundamental stages of the buying process.

- Need Identification:-The buyer becomes aware of an unmet need through advertisement of a product or through friends.
- Product Brokering:-Evaluation of different products to determine which one meets the need.
- Merchant brokering:-Evaluation of merchants based on criteria like price, warranty, etc specified by the buyer.
- Negotiation:-Settling on the terms of the transaction.
- Purchase and delivery.
- Product service and evaluation.

These stages are approximate and may overlap. Software agents can be used in the various stages of the buying process. Agents are suited as mediators in ecommerce due to their personalized, continuous and autonomous nature. The use of agent technologies in the various stages of the buying process helps combat information overload, expedites several stages of the buying process, and reduces transaction costs [33]. Bargain Finder [30] is an agent that helps a buyer in finding a CD vendor who is offering the lowest price; however, the agent is not autonomous in making the purchase decision and a human is still involved in that stage. In most of the current or "first generation" ecommerce applications, buyers are humans browsing through an online catalog of products and making purchases through credit cards. With advances in agent technology, the degree of sophistication and automation is increasing and in the near future automation is expected on both the buyer and seller sides [25]. In "second generation" agent mediated ecommerce, agents will be acting on behalf of buyers and sellers [25]. On the buyer's side, the buyer agent can be expected to help the buyer in identifying desired seller and even negotiate with the seller's agent to maximize the buyer's gains.

In traditional commerce, humans do the thinking; they have the knowledge, are able to deal with uncertainty, and are adaptive by learning from past experiences and by utilizing the knowledge from their trusted sources. In agent mediated ecommerce, agents have to be equipped with these abilities to be able to perform the same jobs as humans do in traditional commerce. The design of agents that are autonomous, intelligent, can deal with uncertainty and incomplete information, and can learn and adapt to a changing environment would be a significant contribution in the area of ecommerce research.

1.2 Overview and Main Contributions

Previous research in the area of buyer agent strategies for choosing seller agents in ecommerce markets has focused on frequent purchases. The work in this dissertation advances the state of the art by presenting reputation based buyer strategies for choosing high quality, low priced sellers which works for frequent as well as infrequent purchases.

This work considers decentralized, open, dynamic, uncertain and untrusted electronic markets where selling agents acting on behalf of sellers are offering goods for purchase to buying agents representing buyers. The goal for a buyer agent (hereafter referred to as *buyer*) is to purchase a product from a seller agent (hereafter referred to as *seller*) who meets its expectations of quality and service, and to purchase it at the lowest price possible in the market. At the same time, buyer wants to reduce its chances of interacting with dishonest and poor quality sellers. In an open market, sellers and buyers can enter and leave the market anytime. In a dynamic market, players in the market need not exhibit the same behavior all the time; sellers can vary the price and the quality in various transactions. Untrusted market implies there could be dishonest sellers in the market. In an uncertain market, buyers can gauge the quality of the product only after actually receiving the product. There could be a onetime transaction between the buyer and the seller or multiple transactions between them. There is no limitation on the number of sellers and buyers in the market. These characteristics are typical of a traditional commerce market and hence an electronic market with similar characteristics is considered in this research.

Products with the following characteristics are considered for sale in the above mentioned electronic market: the quality of goods can be ascertained only after actually seeing it, the quality of the goods varies across the sellers, and the same product by the same seller may be perceived differently by separate buyers depending on their preferences and budgetary constraints. The above mentioned types of goods are classified as "look and feel with variable quality" goods in contrast to commodity products where the quality of the product can be ascertained from afar [13]. Fresh produce, artwork, food items like bread and cakes from specialty stores, and online flower vendors are examples of "look and feel with variable quality" goods, the kind of products we consider sellers to be offering for purchase. Products like books, toys, CD's sold on websites like Amazon are considered quasi commodity products [13] as there is some differentiation within a product niche. A buyer first identifies a particular product to be purchased; the buyer then chooses an online seller based on the price and the reliability of the seller.

It is not possible to pre-program an agent to operate in an open, dynamic, and untrusted market conditions, or to know beforehand who the best seller for a buyer is, as new sellers are entering the market, lowest priced seller may not necessarily be the best seller, and sellers could be lying. Agents have to be equipped with abilities to make the most rational decision based on all the information that they can gather. They should be able to learn from their past experiences.

Recent research has developed intelligent agents for ecommerce applications [8, 9, 14, 18, 30, 49, 50, and 53]. However, as Tran [49] summarizes, agents in [14, 30] are not autonomous, agents in [8, 9, 14, 18] do not have learning abilities, agents in [53] have significant computational costs, and agents in [8, 9, 14, 18, 30, 53] do not have the ability to deal with deceptive agents. Tran and Cohen's work [49, 50] addressed these shortcomings by developing a strategy for buying agents by using reinforcement learning and reputation modeling of sellers. However, their model builds reputation slowly and a buyer has to interact with a seller several times before the seller is considered reputable. This model works well where a buyer has to make repeated transactions with the sellers during frequent purchases. The performance of this model deteriorates for infrequent purchases as the buyer has to purchase several times from a seller before making its decision about the seller. When the buyer is purchasing a product on an infrequent basis, it needs to quickly identify reputed sellers. *First in this dissertation, a reputation based modeling of a seller by the buyer based on its own interactions with sellers which can work for frequent as well as infrequent purchases is presented in a B2C ecommerce market.*

Reputation based adaptive buying agents that reason about sellers for purchase decisions initially explore the market by interacting with various sellers and gradually learn to identify desirable sellers based on their own experiences. During the exploration phase they may interact with undesirable sellers who are charging high prices or offering poor quality or are dishonest and hence they may incur losses. When buyer agents collaborate with honest and similar thinking buyer agent friends by exchanging seller reputation information, they are able to reduce the exploration time and hence the cost associated with exploring [3]. In this dissertation, buyers are considered to be members of online communities in which individuals can join and leave as they wish and can exchange reputation information about sellers with each other in order to identify a trustworthy seller who meets their product expectations and is selling the product at the lowest price.

When buyer agents exchange seller reputation information there is no guarantee that the party replying is honest or, if honest, that they share same opinions regarding product expectations as the buyer agent requesting the information. Opinions can be subjective; consider two sellers *s1* and *s2* selling the same product but at different qualities and prices, and two buyers *b1* and *b2* who purchase the product. Let *s1* be providing mediocre quality at a low price, and *s2* be offering a higher quality product at a higher price. If buyer *b1* is conscious about keeping the costs low and is satisfied with mediocre quality, *b1* will consider *s2* to be expensive and rank seller *s1* higher that *s2*. For buyer *b2*, if high quality is of utmost importance and cost is secondary, *b2* would prefer to buy from seller *s2*. These two buyers cannot use seller information from each other in decision making even if information is exchanged honestly, as their preferences and opinions differ.

Now assume that, buyers *b1* and *b2* share similar opinions and both consider *s2* to be a good seller; however, they may not have the same rating systems. For example, *b1*'s reputation rating for *s2* might be 0.8 on a scale of 0-1 and *b2*'s reputation rating for *s2* might be 80 on a scale of 0-100. When these buyers exchange reputation information, they need to adjust the reputation ratings provided by each other to translate them into meaningful values on their own rating scales. In the previous example, when *b1* obtains *s2*'s rating from *b2*, it has to scale it down 100 times before using the reputation value in its decisions. Hence, before utilizing the information provided by another buyer friend, a buyer has to ensure that they share similar opinions regarding products purchased and the numerical seller rating exchanged be adjusted to account for the differences in the rating systems.

In this dissertation we present two methods of identifying trustworthy buyer friends, who are honest and share similar opinions, and utilizing the information provided by them to choose high quality, low priced sellers.

Specifically, in an open, dynamic, decentralized, uncertain and untrusted market where the same product may be perceived differently by different buyers, the main contributions of this dissertation are:

1. A new reputation based modeling of a seller by the buyer based on its own interactions with sellers, for both frequent and infrequent purchases.

2. A protocol to exchange reputation information about sellers with other buyer friends based on their direct interaction with sellers.

3. Two methods of identifying trustworthy buyer friends who are honest and share similar opinions, and utilizing the information provided by them to maximize a buyer's chances of choosing a high quality, low priced seller.

Chapter 2 Related Work

Chapter 2 presents a review of the relevant work done in related areas. We begin by reviewing the definition of agents and ecommerce followed by a discussion of agents and agent models in ecommerce. In the last section we talk about various trust and reputation models that have been developed.

2.1 Agents

Woolridge [54] defines agents as: "*... a computer system that is situated in some environment and that is capable of autonomous action in this environment in order to meet its design objectives*". Pattie Maes [32] also gives a similar definition "*Autonomous agents are computational systems that inhabit some complex dynamic environment, sense and act autonomously in this environment, and by doing so realize a set of goals or tasks for which they are designed.*" An agent is considered to be intelligent if it is capable of responding to its environment (reactivity), has the ability to take the initiative (proactiveness) and can interact with other agents or humans (social ability) in order to achieve its goals [55]. They may also be adaptive, flexible, temporally continuous, mobile, and communicative with other agents. All software agents are programs but the reverse is not true. For this dissertation we consider software agents as computer programs that exhibit the qualities of agents. They are autonomous, reactive, proactive, and adaptive. They may also be temporally continuous and communicative with other agents.

2.2 Ecommerce

The Electronic Commerce Association defines *ecommerce* [48] as "*electronic commerce [that] covers any form of business or administrative transaction or information exchange that is executed using any information and communications technology.*" It is a very broad definition covering activities like purchasing goods over the telephone, using credit cards to make purchases, automatic billing, order tracking, etc. Commonly, ecommerce refers to buying, selling and exchanging of goods and services over a computer network like the internet. There are four main categories. In business to consumer (B2C) ecommerce, businesses sell to the general public. Online retailing is an example of B2C ecommerce. In business to business ecommerce (B2B), companies do business with each other. For example manufactures or wholesalers may be selling to retailers. In consumer to business ecommerce (C2B), a consumer will make his/her requirements known and businesses will bid to work for the consumer. In consumer to consumer ecommerce (C2C), consumers will be buying and selling to each other. The model presented in this dissertation is meant for a business to consumer ecommerce market.

2.3 Agents in Ecommerce

Guttman and Maes [22] and He et al. [25] present a review of agents that are currently used, the areas in which they can be used and their benefits in ecommerce. According to the Consumer Buying Behavior (CBB) model which covers Business to

Consumer (B2C) transactions [33], there are six fundamental stages of the buying process.

- Need Identification: - Buyer becomes aware of unmet need through advertisement of a product or through friends.
- Product Brokering: - Evaluation of different products to determine which one meets the need.
- Merchant brokering: - Evaluation of merchants based on criteria like price, warranty, etc specified by the buyer.
- Negotiation: - Settling on the terms of the transaction.
- Purchase and delivery.
- Product service and evaluation.

These stages are approximate and may overlap. Software agents can be used in various stages of the buying process. Agents are suited as mediators in ecommerce due to their personalized, continuous and autonomous nature. They can help in the need identifying stage, like detecting when supplies are running low. A notification agent, who has the user profile, can notify the user when a good or a service that matches the user's profile becomes available. For example, in Amazon.com, the latest reviews of titles in categories that match the user's interest are sent to the user. By using agents in the product and merchant brokering stage, search costs of the consumer can be reduced. Using techniques like feature based, constraint, and collaborative filtering, products that match the needs of a user can be recommended.

In the merchant brokering stage, agents can help in finding the appropriate merchant to buy the product from. For example, in BargainFinder [30], if a customer wants to buy a music CD, the agent searches a predefined list of CD sellers to select the CD with the lowest price. Price may not be the only attribute for the user; other attributes can be warranty, delivery and gift services. Shopbot [14] is a comparison price shopping agent. Given the URL of a vendor site, it learns to extract information during the learning period. In the comparison phase, based on knowledge gained during the learning phase, it is able to shop for prices at different vendor sites.

Negotiation is the process of settling for terms and conditions under which the desired product will be purchased and delivered to the customer. Most business to business transactions involve negotiation while most retail sales are fixed price. Automated negotiations can be auctions or bilateral negotiations. In an auction, there is an auctioneer agent who initiates the auction, and bidder agents then make bids according to the bidding protocol. The outcome of the auction is a deal between the auctioneer and the successful bidder. In the digital world many of the impediments that exist in real world auctions (geographical co-location, time expenditure, frustration for the average consumer) disappear. In a bilateral negotiation, the supplier and consumer come to an agreement over the terms of the trade, and the contract is usually multi attributed covering price, quantity delivery date, etc. There is no fixed negotiation strategy; different strategies work in different contexts. He et al. classify the existing work on negotiation in three categories [25]:

- Making decisions by explicitly reasoning about the opponent's behavior.
- Making decisions by finding the current best solution.
- Making decisions by arguing over and above the basic terms of the contract.

Use of agent technologies in various stages of the buying process helps combat information overload, expedite several stages of buying process, and reduce transaction costs.

In the context of CBB model, the proposal presented in this paper encompasses the stages from merchant brokering to service and evaluation. The model proposed uses the reputation of sellers along with the prices quoted to choose a seller in an open, uncertain and dynamic market. The goal is to maximize the buyer's profit and avoid dishonest sellers. The negotiation strategy utilized is to find the current best solution. A seller is evaluated on the product quality and assigned a reputation rating. Reputation of sellers based on direct experiences is exchanged with friends in order to avoid dishonest sellers.

2.3.1 Kasbah

Kasbah [8] was one of the first models of ecommerce marketplace where buying and selling was automated. Users created buying and selling agents which autonomously negotiated and tried to make the best deal on their behalf. The marketplace matched the sellers with the buyers. A user created a selling agent giving it the description of the item to sell, the desired price to sell the item, the desired date by which to sell the item, the lowest acceptable price for the item, and the decay function to reduce the price (linear, quadratic or cubic) for the item. User could also specify whether user approval was needed before finalizing the deal, and whether to send an email notification when the deal was completed. The buying agent was created with the description of the item to buy, the date by which to buy the item, the desired price to pay for the item, the highest acceptable price to pay for the item, and the function to raise the price (linear, quadratic or cubic) for the item. Similar to the selling agent, it could be specified whether user approval was needed before finalizing the deal, and whether to send the user an email notification.

The goal for the buyer agent was to find a seller who was offering the product at the lowest price within a time limit. Seller agents tried to sell their products at the highest price they could sell within a time limit, and they reduced the price of their product with passage of time. Buyer and seller agents autonomously negotiated with each other and tried to make the best deal on behalf of their user. Once transaction was approved, humans took over to transfer cash and goods.

The main advantage of Kasbah was that agents were autonomous; they relieved the user from the process of finding sellers and negotiating with them. The authors pointed out that people are shy of talking to strangers, and their model provided a way for users to avoid it. A buyer agent when choosing a seller agent just focused on the price quoted by the selling agent. It did not consider the quality the selling agent may have provided them in the past. In this model, agents did not learn from their experiences, as the seller agent's previous history with the buyer agent or the reputation of the seller agent was not considered. If the seller agents or buyer agents were dishonest, this model could not identify them. Dishonest agents could be chosen again for interaction in the next episode, if they were offering a price that maximized the utility for the buyer agent or the seller agent.

2.3.2 Equlibria Strategies

Goldman et al. develop strategies for buyers and sellers in electronic markets where there are stock shortages [18]. They consider a market where all sellers sell the same type of good, where each seller sells the same quality product at the same price, and each seller has same units of the product. Buyers and sellers engage in a sequence of encounters. At the beginning of each encounter, the size of stock for each seller is identical. Each buyer is associated with a purchase order and a type which is the average amount of goods that it attempts to purchase in each encounter. A buyer submits its purchase order to only one of the sellers. A seller can sell part, all, or nothing of the amount in the purchase order. Seller's utility increases with the overall units sold and decreases with the units left unsold. Buyer's utility increases with the number of purchase orders satisfied. Seller can buy information about the buyer's type. There are two types of buyers: recognizable buyers whose type information can be used in future encounters, and unrecognizable buyers whose type information can be used only in the encounter it was obtained. Each buyer is associated with only one purchase order and it is valid only for that encounter. The authors develop a representative set of strategies for buyers and sellers and study which profiles are in experimental equilibrium. A profile is a sequence of strategies, one for each buyer and one for each seller. Seller strategies are based on arrival time of the order, size of the order and type of buyer, if the information is available to the seller. Buyer strategies for choosing a seller are based on satisfaction with a seller in the previous encounter.

The strategies developed in [18] are suitable for markets where the buyer is assured of the product quality and fairness in price when choosing any seller in the market, and the buyer's only concern is to get its entire purchase order fulfilled. Our focus is to develop buyer strategies for electronic markets where price and quality vary across sellers, sellers may be dishonest, behavior of sellers may change any time, and seller quotes a price to the buyer only if it can supply the product to the buyer.

2.3.3 Recursive Agent Model

Vidal and Durfeee [53] investigated when agents should act strategically by modeling other agents, or act as simple price takers in an economic multi-agent system. They consider an economy like the University of Michigan Digital Library, where information goods are bought and sold. Because the information economy is rooted in a delivery medium such as the internet, it is virtually free to reproduce the information, there is no shortage of supply, and all agents have direct access to all other agents. The authors point out that agent can survive in such an economy by providing value added services to meet customer's demands.

Their economic model consists of seller and buyer agents. When a buyer agent wants to buy a good, it advertises by requesting for bids. Sellers who have the goods send in their bid. Buyer selects a seller and purchases the good at price p. Buyer can evaluate the quality of a good only after receiving it. Buyer has a valuation function v for each good it wishes to buy; v is a function of price and quality and returns the true value of the product purchased. Buyer's goal is to maximize its product valuation for the transaction. Seller has a selling price and a cost price for a good. The difference of the two is the profit for the seller. Seller's goal is to maximize its profit.

Agents are divided into different classes based on their modeling capabilities. 0-level agents base their actions on inputs and rewards received, and are not aware that other agents are out there. 1-level agents are aware that there are other agents out there, and they make their predictions based on the previous actions of other agents. 2-level agents model the beliefs and intentions of other agents. 0-level agents use reinforcement learning. Agents get some input, take an action and receive a reward. With some probability they explore instead of exploiting the marketplace, and this probability is initially set to 1 and reduced to a minimum value later on. Buyer has a function f for each good that returns the value that the buyer expects to get by purchasing the good at price p. This expected value function is learnt using reinforcement learning as $f = f+\alpha(v-f)$ where α is the learning rate, initially set to 1 and reduced slowly to minimum value. Buyer picks a seller that maximizes its expected value function f. Seller has a function h that returns the expected profit by selling a good at a price p. Seller picks a price that will yield the maximum expected profit h. p is the actual profit and it is the difference between the selling price and the cost price if a bid is made, and 0 if no bid is made. Function h is learned by reinforcement learning as $h = h+\alpha(p-h)$.

1-level buyer stores the history of the last n product qualities sold by seller s. It has a probability density function to calculate the quality to expect from seller s for good g. It uses this quality calculated in computing its expected value function. 1-level seller remembers the last n prices where its bid was accepted by a buyer. It has a probability function to compute the average price that will be accepted by the seller. Seller also maintains a history of what prices were bid by other sellers. It uses both of these to decide at what price to bid.

2-level buyers do not keep deeper models of other buyers or sellers. 2-level sellers model other sellers as 1-level sellers. 2-level sellers first consider what bid the other sellers will make. From these they model which bid the buyer will choose. They then consider what bids they can make to win against the best bid. From their possible bid prices they pick a bid that will yield the highest profit for them.

The main problem addressed in this work is to determine when the agent benefits from having deeper models of others. Their conclusions from the experiments were that, in general, agents having deeper models did better. However, associated with deeper models is the computational complexity and there should be a level at which the gains and the costs of having deeper models balance out for each agent.

Our market model is meant for B2C ecommerce market and considers the existence of dishonest sellers in the market. Buyers use the reputation of sellers to avoid dishonest sellers and reduce their risks of purchasing low quality goods. Reputation of sellers is learnt based on direct interactions and also from information provided by trusted friends. Who can be trusted is learnt, over a period of time, before using the information provided by them in making judgments about sellers. In terms of their 0-1-2 agent model, our buyer is a 1-level buyer, as it makes its decisions based on past behavior of other agents.

2.4 Trust and Reputation Models

When agents do not have complete information about their environment or their interaction partners, there is uncertainty. Trust helps in reducing uncertainty in interactions in open distributed systems. Ramchurn et. al. [37] examine the role of trust in multi agent systems. They conceptualize trust in two ways; system level and individual level trust. System level trust forces agents in the system to be trustworthy because of protocols and mechanisms that regulate the system. Individual level trust models give agents the ability to reason about the trustworthiness of their interaction partners. They state that system level trust can be enforced by designing interaction protocols that force participants to be truthful [44], or develop reputation mechanisms

[12, 29, and 61] that encourage interaction partners to behave in a trustworthy manner, or have security mechanisms [19, 35] that guarantee that new participants are trustworthy.

In [43], Sabater and Sierra provide a review of several computational trust and reputational models. They classify trust and reputation models on the following dimensions.

1. Conceptual Model: Based on their conceptual model, trust and reputation models are classified as cognitive models or game theoretical models. In cognitive models trust and reputation are considered to be the result of an agent's mental state in a cognitive sense [7]. Esfiandiari and Chandrashakeran [15] point out that in cognitive models "*trust and reputation are made up of underlying beliefs and are a function of the degree of these beliefs*." In game theoretical models trust and reputation models are the result of pragmatic games with utility functions, and numerical aggregation of past interactions [1, 6, and 34].

2. Information sources: Trust and reputation models are classified based on the information sources used to compute trust and reputation. Four types of information sources are considered:

- Direct Experiences: There are two types of direct experiences. An agent's own experience with another agent [34, 42, and 61] and an agent's direct observation of other agents interactions [45, 46].

- Witness Information: Information is provided to the agent by other members of the community [6, 61].

- Sociological Information: This is based on the social relations between the agents, and the agents' roles in the society [41, 42]. Examples of social relations can be trade, competition, and collaboration. The roles of agents and their relationships with other agents influence their behavior and their interaction with other agents.

- Prejudice: This is based on assigning reputation to an individual based on signs that identify the individual as a member of a group. Work described in [15, 42] uses this source of information in trust and reputation models.

3. Visibility types: Models are classified based on whether trust and reputation is considered a global property that is shared by all individuals, or a subjective property that is computed separately by each individual. In the first case, trust or reputation of an individual is updated at a central place. It is computed from the opinions of all individuals that have interacted with the individual that is being evaluated. When reputation is treated as a global property, there

is a lack of personalization, since it assumes a common way of thinking. In [6, 56, 57, and 61] reputation is updated at a central place. In the second case, each individual assesses another individual separately based on its own interactions, or based on information gathered from friends or other sources. The work described in [1, 7, 34, 42, 45, 46, 59, and 61] considers trust and reputation as a subjective property.

4. Model's granularity: In this dimension, models are classified based on whether trust and reputation is considered as context dependant or as non context dependant. In a non-context dependant model the trust/reputation value for an individual is a single value [6, 45, 46, and 61]. In a context dependant model the trust/reputation for an individual has different values, one for each context [1, 7, 15, 34, and 42]. For example in a non-context dependant model, the reputation of a seller will be a single value. In a context dependant model, the reputation of a seller can have different values for different contexts like price, quality and delivery date.

5. Agent behavior assumption: There are three different levels of models based on the agent's capacity to deal with cheating agents.

    - Level 0: The model does not consider cheating behavior [6, 59, 15, and 61].

- Level 1: The model assumes that agents can hide or bias the information, but not lie [45].

- Level 2: This model assumes that agents can lie and have mechanisms to deal with it [42, 46].

6. Type of exchanged information: Models are classified on whether they assume boolean information, or continuous measures. Models that use probabilistic methods work with boolean information [1, 46], while models using aggregation mechanisms use continuous measures [6, 34, 42, 45].

7. Trust/reputation reliability measure: Models are classified based on whether the trust/reputation models provide the reliability of the trust/reputation values being computed. Number of experiences, reliability of the witnesses, and age of the information are used to compute the trust/reputation reliability measure [5, 42, and 61].

Our model uses aspects of classifying trust and reputation models, and is conceptually a game theoretical model, that uses direct experiences and witness information, considers reputation/trust as a subjective property, considers reputation/trust to be a single context as it is specifically designed for ecommerce market, assumes that agents can cheat, and information exchanged is a continuous measure. Our model does not compute a separate value for the reliability of the trust value, but instead incorporates the trustworthiness of the witnesses, the experience of the witness, and the age of the information while computing trust and reputation value.

2.4.1 Reputation-Oriented Reinforcement Learning Agent

Tran [49], and Tran and Cohen [50] have developed learning algorithms for buying and selling agents in an open, dynamic, uncertain and untrusted economic market. They used Vidal and Durfee's [53] 0-level buying and selling agents which use reinforcement learning to maximize their utilities as described in Section 2.3.3. They enhance buying agents with reputation modeling capabilities, where buyers model the reputation of sellers. The reputation value varies from -1 to 1. A seller is considered reputable if the reputation is above a threshold value and considered disreputable if the reputation value falls below another threshold value. Sellers with reputation values in between the two thresholds are considered neither reputable nor disreputable. The buyer chooses to purchase from a seller from the list of reputable sellers. If no reputable sellers are available, then a seller from the list of non disreputable sellers is chosen. Initially a seller's reputation is set to 0. The seller's reputation is updated based on whether the seller meets the demanded product value. If the seller meets or exceeds the demanded product value then the seller is considered cooperative and its reputation is incremented. If the seller fails to meet the demanded product value then the seller is considered uncooperative and its reputation is decremented. The 0-level sellers in [20] are augmented with quality varying capabilities along with the price. The seller optionally varies the quality of the product based on its success rate at selling a product to a buyer in order to maximize its profits.

This model builds reputation slowly. A buyer has to interact with a seller several times before the reputation of the seller crosses the threshold value. As this model learns slowly, it is suitable to be used by buyers who are making frequent purchases. In our experiments we show that a buyer using our model for seller evaluation based on direct interactions has a superior performance than agents employing Tran and Cohen's [49, 50] or Vidal and Durfee's [53] models for frequent and infrequent purchases. For frequent purchases, a buyer using our model learns much faster than agents employing Tran and Cohen's or Vidal and Durfee's model, and for infrequent purchases a buyer using our model has significantly higher gains than the other two models.

2.4.2 Regan's Model of indirect assessment of reputation from friends

Regan et al., and Regan and Cohen in [38, 39] provide a model for assessing reputation information provided by friends in an environment where friends could provide deceptive information. The model attempts to account for buyer subjectivity in opinions by partitioning friends into reputable, disreputable and neutral sets. They extend the model proposed by Tran [49] to provide a system where the buyer can request for reputation information about sellers from friends. Buyers from whom reputation information about sellers is requested are termed as advisors. The reputation of advisors is maintained and they are partitioned into reputable, disreputable and neutral advisors. When there are no reputable sellers to consider for purchasing, the buyer requests reputable and neutral advisors for ratings of sellers it is unsure of and for the ratings of some sellers whose reputation is already known. It uses the reputation ratings provided for known sellers to detect if there is a systematic difference between the advisor's rating and its own rating. The mean and the standard deviation of the difference in the ratings is computed and if the standard deviation is low then the advisor's reputation rating for a seller is adjusted by the mean difference to produce a rating for that seller by that particular advisor. All advisors for whom there is a systematic difference in the rating are identified and their ratings are adjusted in the manner described above. An overall rating for a seller is computed from the adjusted ratings of advisors for whom there is a systematic difference. A set of potential sellers is formed, from whom a seller is chosen as described in the Tran model [49]. After purchasing from a seller, the reputation ratings of the advisors are adjusted based on their predictions about the sellers. If the opinion of the advisor and the buyer match, then the reputation of the advisor is increased. If the opinions are in contrast then the reputation of the advisor is decremented. The authors do not specify how the opinion of the friend regarding a seller is inferred from the numerical rating of that seller, nor have they presented any experimental results to justify their theoretical approach.

Our model proposes two approaches of identifying trustworthy friends. In one approach, the trusted friends' opinions and the rating of sellers are utilized along with the buyer's own information to select a seller to purchase from. In the second approach, only the trusted friends' opinions are utilized along with the buyer's own information to select a seller to purchase from. Our work shares some ideas with [38] and [39], in the extent that reinforcement learning is used to adapt a friend's rating and in one of the approaches the mean of the differences between the friend's rating and the buyer's rating of a seller is used to adjust the friend's rating for that seller. Our differences, though, are many more than the similarities: for example, we collect much more complex information from our buyer friends (opinion of the friends about the sellers, reputation value of the sellers, and the experiences of the friends with the sellers), we request friends' suggestion at every transaction, a friend's advice is weighted by its experience, we look at the standard deviation of the rating differences to assure a consistent rating system by our friend, we combine an individual buyer's rating with that of its friends so we can deal with few purchases, and our approaches request for recent information by specifying that friends provide information, only if their interaction with the seller has been after the time specified.

2.4.3 Social Mechanism of Reputation Management

Yu and Singh [59] propose a social mechanism of reputation management for electronic communities. Agent *A* assigns a rating $T_A(B)$ to agent *B* based on, (1) its direct interactions, and, (2) ratings of *B* given by *B's* neighbors and *A's* rating's of those neighbors. The trust rating varies from -1 to 1 and the initial rating is 0. The rating of agent *B* based on direct interactions is increased if it cooperates, and is decreased if it does not. A has two threshold values, w_A and Ω_A, where $-1 < w_A < \Omega_A < 1$. If the rating of B *is* $\geq \Omega_A$, then A trusts B, and if the rating of B *is* $\leq w_A$, than A does not trust B. Each agent has a set of potentially changing neighbors with whom it may interact directly. How an agent determines the reputation of another agent will depend on the testimonies of the latter's neighbors. If $X= (A_0\ldots, A_{N-1}, A_N)$ is a referral chain from agent A_0 to A_N, then A_0 uses the referral chain to compute its trust rating of N at time t ($T_0^X (N)^t$) of A_N as:

$$T_0^X(N)^t = T_o^X(1)^t \otimes \ldots..\otimes T_{N-1}^X(N)^t$$

Where $\otimes$ is the trust propagation factor and is given by:

$$x \otimes y = \begin{cases} x \times y \; if \; x \geq 0 \; and \; y \geq 0, \\ -|x \times y| otherwise \end{cases}$$

Information provided by agent N about B is weighted by the trust rating of N (which is computed through trust propagating mechanism) and is known as testimony E over referral chain and is computed as:

$$E_A^X(N)^t = T_A^X(N)^t T_N^X(B)^t$$

The testimonies over different distinct referral chains are averaged and this value is then used to modify the trust rating of B by A that has been computed based on direct interactions. There could be more than one referral chain from A_0 to A_N and computing trust rating in the above manner using different referral chains may result in different and possibly conflicting values.

In this model subjectivity in opinions and possibility in variations in trust rating systems is not addressed. In our model, buyers are members of online communities like online forums, blogs, and discussion boards. Buyers request information on sellers from other members of the online community. As the buyer does not know other members, who respond personally, it cannot assume that they are honest, or share similar opinions, as opinions can be subjective. For buyers with similar opinions there is also no guarantee that they share the same rating systems. Our work addresses issues of subjectivity in opinions and variations in the seller rating systems of different buyers by providing a method of identifying trustworthy friends (honest and of similar opinions) from all the members of community who respond with information on sellers.

2.4.4 REGRET and FIRE

REGRET [41, 42] is a model that takes individual dimension (direct interactions), social dimension, and hierarchical ontology structure of agents into account.

Individual Dimension: - This is based on direct interaction between two agents. If two agents A and B interact, the rating of B at time t as computed by agent A on an aspect (price, quality, etc.) is the weighted mean of the individual ratings, giving more importance to recent ratings.

Social Dimension- According to REGRET, an individual by default inherits the reputation of the group it belongs to. Also, an individual will use the experiences of the members of his/her group to complement his/her experiences. Here they consider three types of social reputation ratings, one based on personal experience and two based on social experience.

- The reputation rating based on the personal experience of agent *A* with the members of the group to which agent *B* belongs on a particular aspect.

- The reputation rating based on what the members of the group to which agent *A* belongs think about *B* on a particular aspect.

- The reputation rating based what the members of group to which agent *A* belongs think about the group to which agent *B* belongs on a certain aspect.

The reputation rating combining the individual reputation rating and the three social reputation ratings is defined as the weighted sum of the individual components, where the weight of each component is set by the user.

Ontological Dimension: - The reputation rating computed above is based on a single aspect. In REGRET, reputation is considered as a multifaceted concept. The ontological dimension is the combination of reputation on different aspects to produce a single rating. Each agent uses an ontological structure to combine the ratings, and the importance of each individual rating in producing the combined rating is set by the agent based on its preferences. For example, the reputation rating of agent *B* as a good seller as computed by agent *A* can be the weighted sum of the reputation rating of *B* for individual aspects like price, quality, and delivery date.

REGRET gives a general framework for combining individual trust ratings, but does not specify how the individual interaction trust ratings are to be computed. The advantage of this model is that it takes into account the social dimension and the ontological dimension. However, as researchers in [24] point out, the social dimension component is dependant upon the agent's social network and the authors do not specify how to build and expand the social network. The size of the agent's social network determines the number of members involved in computing the social dimension component.

FIRE [23, 24] is a modular trust and reputation model consisting of four main components:

1. Interaction trust is the trust resulting from direct interactions between two agents. The direct trust component of REGRET [41, 42] is used to model this part. For each interaction and term (price, quality, etc.) between two agents, *A* and *B*, *A* assigns a rating for *B*. A history of past ratings is stored in the local database. The trust rating for agent *B* for a particular term is computed as the weighted mean of all the ratings for that particular term in the database. More importance is given to more recent ratings.

2. Role-based trust (RT) models the trust resulting from role-based relationships between two agents. Rules are used to assign RT values. Each rule is a tuple specifying the roles of the agents, the term, the rating if the agents are in that particular role and the default level of the influence of the rule. To compute RT all relevant rules between agents *A* and *B* are looked up. RT is computed

as the sum of the product of the rating of the rule and the influence level of the rule divided by the sum of the influence levels.

3. Witness Reputation (WR) is based on finding witnesses that have interacted with that particular agent. A referral system is used to find such witnesses. The WR rating is computed as the weighted mean of all the ratings provided by witnesses.

4. Certified Reputation (CR) are the ratings presented by the rated agent about itself which have been obtained from agents with whom it has interacted in the past.

The overall rating is computed as the weighted mean of all the individual trust components. The relative importance of each component is set by the end user. Each kind of rating has a reliability rating which reflects the confidence of the trust model in producing that information. Reliability rating is measured from two values: 1) The number of ratings that have been considered to compute a trust rating; and 2) the rating deviation reliability. If the agent's behavior is volatile then its ratings will vary, as reflected by the rating deviation reliability.

Both FIRE and REGRET give a framework for combining individual trust ratings, but do not specify how the individual interaction trust ratings are to be computed, as their trust and reputation systems are designed for general open multi agent systems. Our model is designed for multi agent systems in an ecommerce environment and provides a method for computing the direct rating of the sellers based on the buyer's interaction with the sellers. Our goal is to optimize the buyer's performance by helping it choose a seller who is offering the quality it desires at the lowest price. Various sellers can be offering different qualities of a product at different prices. For honest sellers, we expect price to be proportional to quality. Buyers' preferences may be different. Some may be satisfied with medium quality because of their budgetary constraints while others may want the highest quality irrespective of the cost. In our model, the seller's ability to meet the buyer's expectations about quality (product quality, delivery schedule as expected by the buyer, etc.) and the price at which the product was purchased in relation to price bids made by other sellers is used to compute the buyer's direct rating of the sellers.

In REGRET witnesses are obtained through the agents' social network. The number of witnesses is limited by the size of the agent's social network. In FIRE witnesses are obtained through referral system [59] and assumed to be truthful. As mentioned previously, our model is designed for ecommerce domains like the internet. It is common practice for online human buyers to post queries regarding products or vendors at various user groups to gather more information. We extend the same concept for the buyer agents in our model. To find friends who can provide information on sellers, we consider buyers to be members of online communities like online forums, blogs, and discussion boards. Buyers request other members of the online community for information on sellers. As the buyer does not personally know other members who respond, it cannot assume that they are honest, or share similar opinions, as opinions can be subjective. For buyers with similar opinions there is also no guarantee that they share the same rating system. In our work we provide methods for identifying trustworthy friends (honest and of similar opinions) from all the members of community who respond with information on sellers. In FIRE and REGRET, the issue of subjectivity in opinions and the possibility of variation in rating scales of witnesses are not addressed.

In our work the rating of a trustworthy friend with a larger number of buying experiences is weighted more than a trustworthy friend whose buying experiences are fewer. We do not consider role-based trust as our model is meant for B2C business model where normally there are no other relationships between the buyer and the seller other than buyer and seller. The concepts of role-based trust and certified reputation could be considered in the future extensions of the model.

2.4.5 Histos and Sporas

Histos and Sporas [61] are two centralized mechanisms for maintaining reputations in online communities and marketplaces.

Sporas provides a global reputation value for each member in the online community. In Sporas, new users start with a minimum reputation value and build their reputation during their activities in the system. The reputation value of a user does not fall below the value of a newcomer even if the user is very unreliable. This is to prevent or discourage the user from leaving the system and entering with a new identity. After each transaction the reputation value of the user is updated by the other party to reflect his/her trustworthiness in the last transaction. Two users can rate each other only once. If the users interact more often, then only the last submitted value is kept in the system. This is to prevent two parties from colluding with each other and inflating their reputation values. A user with high reputation values experience smaller change in their ratings after each update. When updating the rating of a user the reputation value of the user giving the rating is also taken into account.

Histos is a more personalized system than Sporas. In Histos different reputation ratings connecting the users of the system are considered. Pair wise ratings in the system are represented as a directed graph, where nodes represent users and weighted edges represent the most recent reputation rating given by one user to another, with the direction pointing toward the rated user. If a connected path exists between two users, A and B, then a more personalized reputation value of B can be computed for A. The reputation of an agent at level L of the graph (L>0) is calculated recursively as a weighted mean of the rating values that agent in level L-1 gave to that agent. The weights are the reputations of the agent rating the target agent. Any agent who has been rated directly by the agent owner of the graph has a reputation value equal to the rating value. This is the base case of the recursion. The model limits the length and the number of the paths that are taken into account for calculation.

As Histos and Sporas are centralized reputation systems in which the reputation of a user is updated by other parties at a central place, it is vulnerable to multi party collusion. There is no guarantee about the accuracy of the rating. By having a single global rating, the subjectivity of the parties rating the user cannot be accounted for.

Our trust and reputation model is decentralized in which each buyer agent separately maintains the reputation of seller agents it has interacted with. Buyer agents exchange seller reputation information with each other when evaluating sellers for purchase decisions. Our model addresses subjectivity in buyers' opinions and the possibility of variations in buyers' ratings of sellers.

Chapter 3 Model

We consider a decentralized electronic market, where agents representing sellers offer products, where agents representing buyers purchase the products, and where the quality of the product can be gauged only after receiving it from a seller. Seller and buyer agents can enter and leave the market at any time, they may be deceptive, they may change their nature at any time, there can be any number of seller and buyer agents, and there can be one time or multiple transactions between them. From this point onwards we refer to buyer and seller agents as *buyers* and *sellers*. The *buyer* has a valuation for the product which is based on the quality and price of the product. The buyer's goal is to purchase from a *seller* who will maximize its valuation of the product. At the same time it wants to avoid interactions with dishonest or poor quality sellers in the market.

The buyer models seller's reputation based on its direct interactions with it. The reputation of a seller reflects the seller's ability to provide the product at the buyer's expectation level, and its price compared to its competitors in the market. It is common practice for online human buyers to post queries regarding products or vendors at various user groups to gather more information. We extend the same concept for buyers in our model. Each buyer in this model belongs to a community comprising of buyers. Members of the community are referred to as *friends* in this model. While evaluating sellers for purchase decisions, the buyer requests for information on sellers from *friends*. There is no guarantee that friends are telling the truth or are using the same evaluation mechanism of sellers as this buyer. The buyer models its friends' reputation to identify *trustworthy friends* who are honest and share similar opinions regarding product qualities. The buyer then calculates an aggregate trust rating for each seller based on an *average rating* of the seller computed from its direct interactions with that seller, and an *average reputation rating* of that seller which is computed based on the information provided by trusted friends. An *aggregate trust rating* is used to identify reputable sellers as potential sellers, and the buyer chooses to purchase from a reputable seller who maximizes its product valuation among all potential sellers.

The following notation is used. Subscript represents the agent computing the rating. Superscript represents the agent about whom the rating is computed. Information in parentheses in the superscript is the kind of rating being computed. For example, every time buyer *b* purchases a product from seller *s*, buyer *b* computes direct trust *(di)* rating $T_b^{s(di)}$ of seller *s*.

An aggregate trust rating T_b^s for each seller *s* is computed from the following components.

- Average Direct Trust Rating $T_b^{s(diavg)}$ - This rating is computed based on direct interactions between buyer *b* and seller *s*.
- Average Reputation-Based Trust Rating $T_b^{s(avgrep)}$ - This rating is computed based on reputation of *s* gathered from *b*'s trustworthy friends.

3.1 Average Direct Trust Rating ($T_b^{s(diavg)}$)

Average direct trust rating $T_b^{s(diavg)}$ for seller s is computed based on buyer b's direct interactions with s. Buyer b records a history of its transactions with sellers as a tuple $(s, t, p, pr, q, T_b^{s(di)})$, where s is the seller's identification, t is the time of the interaction, p is the product, pr is the price, q is the quality, and $T_b^{s(di)}$ is the trust rating of s based on direct interaction and is computed as described below.

Each time buyer b purchases from a seller s, based on the quality of the product received and the price paid to purchase it, the trust rating of seller s by buyer b is computed as shown in equation (1).

$$T_b^{s(di)} = \begin{cases} \dfrac{q_{act}}{q_{exp}} - \left(\dfrac{p_{act} - p_{avg}}{p_{max}}\right) if\ q_{act} \geq q_{min}\ and\ p_{act} \geq p_{avg} & (a) \\ \dfrac{q_{act}}{q_{exp}}\ if\ q_{act} \geq q_{min}\ and\ p_{act} < p_{avg} & (b) \\ \dfrac{q_{act}}{q_{exp}} - \left(\dfrac{p_{act} - p_{min}}{p_{max} - p_{min}}\right) if\ q_{act} < q_{min} & (c) \end{cases} \quad (1)$$

Where q_{act} is the actual quality of the product delivered by seller s, q_{exp} is the highest quality buyer b expects, q_{min} is the minimum quality expected by buyer b, p_{act} is the price paid by buyer b to purchase the product from seller s, p_{min} is the minimum price quote received, p_{max} is the maximum price quote received, and p_{avg} is the average of the price quotes received by buyer b for this product.

Buyer b uses two components to calculate the trust rating of seller s: the degree to which the quality delivered by s is meeting b's expectation, and the price charged by s in comparison to the price quotes that b has received from other sellers. Buyer b's trust rating of a seller *s* is directly proportional to the ratio of the quality *s* delivers, to *b*'s desired quality. If there are two sellers, *s1* and *s2*, who offer same quality, and the price asked by *s1* is lower than *s2*, then buyer *b* would like to give *s1* a higher rating than *s2*. To modify the trust rating of a seller based on price, *b* compares the price charged by seller *s* to all price quotes it has received. The trust rating of seller *s* is decreased based on how high *s*'s price is compared to its competitors.

Since quality is of utmost importance, *b* first categorizes seller *s* based on whether quality delivered by *s* (q_{act}) meets or exceeds *b*'s minimum quality expectation (q_{min}). Equations (1a) and (1b) are used to compute the trust rating of a seller when the buyer's quality expectation criterion is met. Equation (1c) is used to compute the trust rating for a seller in cases where the seller fails to meet the buyer's minimum quality expectation.

Like other researchers [39, 40], we make the common assumption that it costs more to produce a higher quality product. So when seller *s*'s quality is greater than buyer *b*'s minimum quality expectation, *b* compares the price charged by *s* to the average price that was quoted for the product. If *s*'s price is greater than the average price quoted, the difference between *s*'s price and the average price quoted is weighed against the maximum price quoted for that product (part (a) of the equation (1)). If *s*'s price is below the average price (which can happen if other sellers are trying to maximize their profits, or if there are too many low quality sellers), then the rating for *s* is computed based on its quality alone (part (b) of equation (1)).

When seller s fails to deliver b's minimum quality expectation, b wants to give very low rating to s. So, the difference of s's price and the minimum price quoted amongst various sellers is weighted against the difference between the maximum and minimum prices quoted to penalize the seller severely (part (c) of the equation (1)).

This model makes two assumptions:

- When quality delivered by a seller is higher than buyer's highest quality expectation, then buyer regards the seller as meeting its highest quality expectations. In other words, if $q_{act} > q_{exp}$, q_{act} is set as equal to q_{exp}.
- It costs more to produce higher quality products.

$T_b^{s(di)}$ ranges from -1 to 1. In the best case, b gets the expected quality at the lowest price and $T_b^{s(dimax)} = 1$ (equation 1b). In the worst case, $q_{act} = 0$; b pays the maximum price quoted and $T_b^{s(dimin)} = -1$ (equation 1c).

Whenever buyer b is evaluating a list of sellers for purchasing decisions, it computes $T_b^{s(diavg)}$, the average rating for each seller s from its past interactions. $T_b^{s(diavg)}$ is computed as the weighted mean of its past n recent interactions with s.

$$T_b^{s(diavg)} = \frac{\sum_{i=1}^{n} w_i T_{b(i)}^{s(di)}}{\sum_{i=1}^{n} w_i} \tag{2}$$

Where $T_{b(i)}^{s(di)}$ is the rating computed for a direct interaction using equation (1). Subscript i indicates the i^{th} interaction, and w_i is the importance of the rating in computing the average and computed as:

$$w_i = \frac{t_{cur}}{t_{cur} - t_i} \tag{3}$$

Recent ratings should have more importance. Hence the weight of a rating is inversely proportional to the difference between time t_i a transaction happened to the current time t_{cur}. Buyer b has threshold values θ and ω for direct trust ratings to indicate its satisfaction or dissatisfaction with a seller respectively. Threshold values θ and ω are set by the buyer b, $\theta > \omega$, and θ and ω are in the range [-1, 1]. Buyer b partitions the sellers into reputable S_R, disreputable S_{DR} and neutral $S_?$ seller categories. If S is the list of sellers, then the criteria for a seller to belong to one of the categories is:

$$\begin{gathered}
\forall\, s \in S \mid s \in \begin{cases} S_R \quad if\ T_b^{s(diavg)} \geq \theta \\ S_{DR}\ if\ T_b^{s(diavg)} \leq \omega \\ S_?\ otherwise \end{cases}, \\
S = S_R \cup S_{DR} \cup S_?, \\
S_R \cap S_{DR} = 0, S_R \cap S_? = 0, S_{DR} \cap S_? = 0,
\end{gathered} \tag{4}$$

Sellers are considered reputable if their average direct trust rating $T_b^{s(diavg)} \geq \theta$, disreputable if the average trust rating $T_b^{s(diavg)} \leq \omega$, and buyer is unsure about sellers whose average direct trust ratings fall between ω and θ. Buyer b recommends sellers in the reputable list S_R, does not recommend sellers in the disreputable list S_{DR}, and expresses its uncertainty about sellers in the neutral list $S_?$.

3.2 Average Reputation-Based Trust Rating ($T_b^{s(avgrep)}$)

This value is computed from the information provided by *b*'s trustworthy buyer friends who are honest and share similar opinions.

As mentioned previously, each buyer in this model belongs to a community of buyers. Buyer *b* requests for information on sellers from members of the community. Members of the community are referred to as *friends* in this model; however, there is no guarantee that friends are telling the truth, or they are using the same evaluation mechanism of sellers as buyer *b*. Hence it is important for *b* to identify trustworthy friends who are truthful and whose opinions regarding sellers are similar to *b's* opinions. It then has a circle of trustworthy *friends*, that it can use in cases where it has no direct previous experience with a seller. In this dissertation, two methods of identifying trustworthy friends are presented. In both methods, over a period of time based on the opinions sent by *b*'s friends and its own experience with the sellers, *b* learns which of the friends can be trusted in recommending good sellers and compiles a list RF_b of trustworthy friends. For each friend *f* who responds to a request for information by buyer *b* regarding the quality of a seller, *b* models *f*'s reputation r_b^f. The value of r_b^f ranges from -1 to 1 and is initially set to 0. Each time *f* makes a seller recommendation that corresponds to the experiences of *b* with that same seller, its reputation is incremented; every time *f* makes an inaccurate recommendation its reputation is decremented.

3.2.1 First Approach to Identifying Trustworthy Friends

In the first approach the process of identifying trustworthy friends works as follows: Buyer *b* compiles *S?*, a list of sellers, it has received bids from and requests information from friends regarding these sellers. For each seller *s* in the list *S?, b* requests the following information from each friend *f*:

- Friend *f's* opinion, which is one of "recommend", "not recommend" or "neutral".

- w_f^s, number of direct interactions , that *f* has had with seller *s* .

- A numerical rating R_f^s for seller *s*, where R_f^s is the rating by friend *f* for *s*.

Buyer *b* stores the responses from its friends, and after *b* has purchased from a seller *s*, it compiles a rating for that seller as described in the section for trust based on direct interactions. Next, *b* compares the responses of its friends to its own experience. If *b*'s opinion and friend *f*'s opinion regarding *s* match and the reputation of *f* is below a threshold value ψ, then the reputation of *f* is incremented. This happens in two cases:

1) If buyer *b* was happy with its own experience (considers seller *s* to be reputable) and friend *f* also had recommended *s* , or,

2) If *b* was unhappy (considers *s* to be disreputable) and *f*'s opinion was also "not recommend".

The reputation of friend f is incremented by:

$$r_b^f = r_b^f + \alpha\left(1 - r_b^f\right) \, if \, r_b^f \geq 0$$
$$r_b^f = r_b^f + \alpha\left(1 + r_b^f\right) \, if \, r_b^f < 0 \tag{5}$$

If buyer b's and friend f's opinion do not match, then the reputation of friend f is decremented by:

$$r_b^f = r_b^f + \beta\left(1 - r_b^f\right) \, if \, r_b^f \geq 0$$
$$r_b^f = r_b^f + \beta\left(1 + r_b^f\right) \, if \, r_b^f < 0 \tag{6}$$

Where α and β are positive and negative factors respectively and chosen by buyer b.

Buyer b also computes $\Delta_b^{f(s)}$, the difference between friend f's rating and its own rating for seller s, and stores it. $\Delta_b^{f(s)}$ is computed as:

$$\Delta_b^{f(s)} = T_b^{s(diavg)} - R_f^s \tag{7}$$

Where $T_b^{s(diavg)}$ is buyer b's average rating for seller s based on its direct interactions with it, and R_f^s is the friend's f's rating for seller s. To determine if friend f is consistent in its opinions, standard deviation, σ, of m consecutive differences, $\Delta_b^{f(s)}$, between f's and b's ratings for sellers is computed as:

$$\sigma = \sqrt{\frac{1}{m}\sum_{s=1}^{m}\left(\Delta_b^{f(s)} - \Delta_b^{f(avg)}\right)^2} \tag{8}$$

Where m is set by buyer b, and $\Delta_b^{f(avg)}$ is the average of m consecutive differences between the ratings of trusted friend f and buyer b for sellers that the buyer b has transacted with in the last m transactions. $\Delta_b^{f(avg)}$ is computed as:

$$\Delta_b^{f(avg)} = \frac{\sum_{s=1}^{m} \Delta_b^{f(s)}}{m} \tag{9}$$

If the reputation value r_b^f for friend f exceeds the threshold value ψ, and if the standard deviation, σ, of m consecutive differences between f's and b's ratings is small and below μ, then friend f is added to the list of trustworthy friends RF_b. Threshold values ψ and μ are set by the buyer. By ensuring that the standard deviation is small, only a friend who is consistent in its rating mechanism is added to RF_b. On the other hand if the reputation of the friend no longer exceeds the threshold value ψ or the standard deviation increases and is no longer below μ, then the friend is removed from RF_b.

Buyer agent b does not only consider the recommendations of trustworthy friend buyers, but it also adapts the trustworthy friends' rating to better match its own rating algorithm. To do so, $\Delta_b^{f(avg)}$, the average of m consecutive differences between trustworthy friend f's rating and buyer b's rating is used to adjust the rating of a seller from that trustworthy friend f. The number of experiences as reported by trustworthy friend f is used to decide whose opinion should be weighed more (for example, a trustworthy friend whose opinion is based on 30 experiences should be weighed more than that of a trustworthy friend whose opinion is based on a single experience). $T_b^{s(avgrep)}$, the average rating of seller s across trustworthy friends is computed as:

$$T_b^{s(avgrep)} = \frac{\sum_{f=1}^{n} w_f^s \, (\Delta_b^{f(avg)} + R_f^s)}{\sum_{f=1}^{n} w_f^s} \tag{10}$$

Where n is the total number of friends in RF_b , and w_f^s is the number of experiences of trustworthy friend f with seller s, and $\Delta_b^{f(avg)}$ is computed as in equation (9).

3.2.2 Second Approach to Identifying Trustworthy Friends

A second approach that will be examined is when the buyer utilizes only opinions and experiences of trusted friends regarding sellers. In this case, the buyer only requests for friend's opinions and the number of times a friend has interacted with sellers. The reason for studying this approach is to determine whether it enables the buyer to utilize the opinions and experiences of a higher number of friends, the rationale being that there will be a consistent difference between a buyer's rating and a friend's rating only if the rating mechanisms of the buyer and the friend follow the same curve. Friends may be honest and their recommendation for a seller may be similar to the buyer's, however if the difference in ratings is not constant, then their opinions will not be considered in the first approach.

In the second approach, buyer b compiles a list *S?* of sellers it has received bids from and requests information from friends regarding these sellers. For each seller s in the list *S?, b* requests the following information from each friend f:

- Friend *f's* opinion, which is one of "recommend", "not recommend" or "neutral".

- The number of direct interactions w_f^s, that f has had with seller s.

Buyer b stores the responses from its friends, and after b has purchased from a seller s, it compiles a rating for that seller as described in the section for trust based on direct interactions. Next, b compares the responses of its friends to its own experiences. If b's opinion and friend f's opinion match, and the reputation of friend f is below a threshold value ψ, then the reputation of friend f is incremented as given in equation (5). If buyer b's and friend f's opinions do not match, then the reputation of friend f is decremented as given in equation (6).

If the reputation value r_b^f for friend f exceeds the threshold value ψ, then friend f is added to the list of trustworthy friends RF_b. Buyer b assigns a reputation rating R_f^s to seller s based on the recommendation of trusted friend f. R_f^s is assigned a value as shown in equation (11).

$$R_f^s = \begin{cases} \theta, if\ f\ recommends\ s \\ \omega, if\ f\ does\ not\ recommend\ s \\ \gamma, where\ \omega < \gamma < \theta, if\ f\ is\ neutral\ about\ s \end{cases} \quad (11)$$

At this point, buyer b has determined that f is honest and has similar opinions regarding sellers. So, when f recommends s, since b would recommend s only if its trust rating was greater than or equal to θ, b sets $R_f^s = \theta$. Since b does not recommend a seller when its trust rating is less than or equal to ω, b sets $R_f^s = \omega$, when f does not recommend s. When f expresses its uncertainty over s, then R_f^s is set to a value between ω and θ.

The average rating of seller s across friends is computed as:

$$T_b^{s(avgrep)} = \frac{\sum_{f=1}^{n} w_f^s R_f^s}{\sum_{f=1}^{n} w_f^s} \qquad (12)$$

Where n is the total number of trusted friends in RF_b, and w_f^s is the number of experiences of trusted friend f with seller s.

3.3 Trust Rating for an Unknown Seller ($T_b^{s(new)}$)

When buyer b cannot infer a rating for seller s based on its information or based on information provided by trusted friends, it classifies seller s as a *new seller*.

3.4 Aggregate Seller Rating (T_b^s)

After receiving price bids from various sellers, buyer b evaluates sellers based on its own knowledge and information from trusted friends, and computes a net trust rating T_b^s for each seller s. The following scenarios are possible.

Case1: When seller s is new to buyer b ($T_b^{s(diavg)}=0$) and to its trusted friends ($T_b^{s(avgrep)}=0$), then $T_b^s = 0$.

Case 2: When seller s is new to buyer b's trusted friends ($T_b^{s(avgrep)}=0$), but buyer b has previously purchased from seller ($T_b^{s(diavg)}\neq 0$), then $T_b^s = T_b^{s(diavg)}$.

Case 3: When seller s is new to buyer b ($T_b^{s(diavg)}=0$) but at least one of its trusted friends has purchased from seller s in the past ($T_b^{s(avgrep)} \neq 0$), then $T_b^s = T_b^{s(avgrep)}$.

Case 4: When seller s is not new to buyer b or its trusted friends ($T_b^{s(diavg)} \neq 0$) and ($T_b^{s(avgrep)} \neq 0$), then:

$$T_b^s = xT_b^{s(diavg)} + yT_b^{s(avgrep)} \tag{13}$$

Where $x+y=1$ and the values of x and y are set by buyer b to reflect how much weight is to be given to each component when both are available.

3.5 Seller Selection

When buyer b wants to purchase a product, it requests price bids from sellers who are interested in selling the product. After obtaining the price bids, buyer b then requests information from all its friends on all the sellers who have sent price bids. However, it only utilizes information about sellers provided by *trustworthy friends* who are in its trustworthy friends list RF_b. A seller is classified as new if neither the buyer nor any of its friends in the RF_b have interacted with that seller previously.

After computing the aggregate trust rating T_b^s for each seller s, buyer b identifies a set of potential sellers who's T_b^s is above the satisfaction threshold θ. Whether new sellers should be added to the list of potential sellers depends on the attitude of the buyer. We consider two attitudes:

1. *Conservative*:-This kind of buyer prefers to deal with known and reputable sellers, and is only willing to try out new sellers if it has not yet found good reputable sellers. It explores the market more frequently in the beginning and occasionally at a later stage. From the list of sellers who have sent bids, the buyer forms a list of sellers about whom it could not infer information based

on its own and its trusted friends' experiences. These new sellers are considered as *unexplored*, and the exploration rate is set proportional to the ratio of unexplored sellers to all the sellers who have sent bids. As the number of unexplored sellers decreases, the exploration rate also comes down and is prevented from further reduction after it has reached a minimum rate.

2. *Risk Taking*:- This kind of agent always includes new sellers into the list of potential sellers to be able to quickly identify a good seller.

The buyer has a valuation function for the product, which is a function of the price a seller is quoting and the quality that has been delivered in the past. For a seller with whom the buyer has interacted before, the quality is the average of the quality delivered in the past interactions. For a seller with whom the buyer has not interacted directly, the quality is set to the desired expected quality. From the list of potential sellers, the buyer chooses a seller who maximizes its product valuation function.

3.5 Summary

Our buyer model is designed for decentralized, open (sellers and buyers can enter and leave the market at any time), uncertain (quality of the product can be gauged only after actually seeing the product) and untrusted (sellers and buyers may be dishonest) electronic markets, where various sellers (honest, dishonest, overpriced, and offering varying quality) offer products for purchase. The buyer's goal is to purchase its desired product from a seller who will maximize its product valuation. The buyer's valuation of the product is based on the quality and price of the product.

Each time the buyer purchases a product from a seller, it computes a trust rating for that seller. It uses two components to calculate the trust rating of that seller for that interaction: the degree to which the quality delivered by seller is meeting buyer's expectation, and, the price charged by the seller in comparison to price quotes the buyer has received from other sellers. The average rating of that seller based on direct interactions is computed as the weighted mean of trust ratings of the buyer's past interactions with the seller where the weight of a trust rating is proportional to its recentness.

Each buyer in this model belongs to a community comprising of buyers who are referred to as *friends.* While evaluating sellers for purchase decisions, the buyer requests for information on sellers from friends. There is no guarantee that friends are telling the truth or are using the same evaluation mechanism of sellers as this buyer. The buyer models its friends' reputation to identify *trustworthy friends* who are honest and share similar opinions regarding product qualities.

In this model we present two methods for a buyer to identify trustworthy friends. In the first method, the buyer utilizes other friends' opinions and ratings of sellers to identify trustworthy buyer friends. The reputation of a seller provided by trustworthy friends is adjusted to account for the differences in the rating systems. The average reputation rating of a seller is computed as the weighted mean of adjusted reputation ratings of that seller provided by trusted friends. The weight of a rating is proportional to the experience of a trusted friend. In the second method, the buyer only utilizes other buyers' opinions of sellers to identify trustworthy friends. Ratings are assigned to sellers by the buyer based on trustworthy friend's opinions. The average reputation rating of a seller is computed as the weighted mean of the assigned ratings.

The buyer *calculates* an aggregate trust rating for each seller based on the average rating of the seller computed from its direct interactions with that seller, and the average reputation rating of that seller computed based on the information provided by trusted friends. The aggregate trust rating is used to identify reputable sellers as potential sellers. Whether new sellers are added to the list of potential sellers depends on the attitude of the buyer. A buyer with risk taking attitude includes new sellers in the list of potential sellers. A buyer with conservative attitude does not include new sellers in the list of potential sellers, and only tries out new sellers when it wants to explore. From the list of potential sellers, the buyer chooses a seller who maximizes its product valuation function.

Chapter 4 Experiments and Results

In this chapter, we describe the experiments conducted to evaluate our model, and we analyze the results. In our model:

- We presented a method for a buyer to evaluate sellers based on its direct interactions.
- We presented two methods of identifying trustworthy (honest and share similar opinions) friends and utilizing information from them.
- We considered risk taking and conservative attitude of buyers.

We tested various components of the model in five phases. In Phase I, our goal was to test buyer's method of evaluating sellers based on direct interactions. In a market populated with various sellers (varying quality, honest, dishonest, and varying prices), we compared the performance of a buyer using our method of rating sellers based on direct interactions with buyers using models, for rating sellers based on direct interactions, proposed by other researchers. The time taken by various buyers to learn to identify high quality, low priced sellers was used as a metric to compare various buyers. We want buying agents to identify high quality sellers offering low prices as soon as possible. If a buyer is able to identify high quality sellers quickly, then the buyer's strategy can be used when making infrequent purchases.

In Phase II, our goal was to test first method of identifying and utilizing information from trustworthy friends as described in section 3.2.1. In this method, the buyer requests for friend's opinions, and ratings of sellers. The buyer uses friend's opinions and standard deviation of the differences between friends' ratings and it's ratings of sellers to identify trustworthy friends. The reputation ratings of sellers provided by trustworthy friends are adjusted to account for the differences in the friends' and the buyer's rating systems and combined with the buyer's rating of sellers to identify potential sellers to purchase from.

In Phase II, the market was populated with various sellers (varying quality, honest, dishonest, and varying prices) and with different types of buyer friends (honest, lying, similar opinions, and different opinions). We analyzed the first method: 1) in successfully identifying trustworthy (honest and sharing similar opinions) friends, 2) utilizing the information (friend's opinion regarding the seller and the numerical rating for that seller) provided by these trustworthy friends along with its own information regarding sellers in making purchase decisions, and, 3) determine, if the first method helps in improving the performance of a buyer with friends as compared to a buyer making decisions based on its own information. The metrics used were: ability to identify trustworthy friends, time taken to identify trustworthy friends, and percentage of purchases from honest, high quality sellers at different levels of purchase, and as the number of sellers in the market was varied. Desired behavior is that honest friends who share similar opinions are identified at the earliest and their information about sellers utilized all the time. Friends who are not honest, or who do not share similar opinions should not be identified as trustworthy and their information regarding sellers should not be utilized at all. Buyer with friends should make higher percentage of purchases than buyer acting alone from honest, high quality sellers.

In Phase III, we evaluate second method of identifying and utilizing information from trustworthy friends as described in section 3.2.2. In this method, buyer only requests for friend's opinions of sellers. Buyer uses friend's opinions to identify trustworthy friends. Based on a trusted friend's opinion, reputation value is assigned to a seller. Assigned reputation values and the buyer's ratings of sellers are combined to identify potential sellers to purchase from. In Phase III, the market was populated with various sellers (honest, dishonest, varying quality, and varying price) and different types of friends (honest, lying, sharing similar opinions, and having different opinions). We compared the performances of buyers with friends using first and second method of identifying trustworthy friends. Metrics were: their abilities to identify trustworthy friends, time taken to identify trustworthy friends, and percentage of purchases from honest, high quality sellers at different levels of purchase, and as the number of sellers in the market was varied. Desired behavior is that honest friends who share similar opinions are identified at the earliest and their information about sellers utilized all the time. Friends who are not honest, or who do not share similar opinions should not be indentified as trustworthy and their information regarding sellers should not be utilized at all. We wanted to determine if any method (first or second) of identifying trusted friends and using information from them is superior.

In our model, we consider two attitudes of buyers: risk taking and conservative. A risk taking buyer considers new sellers (buyer has no information on these sellers) as reputable initially and tends to purchase from a new seller if it is offering the lowest price among reputable sellers. A buyer with conservative attitude is cautious in its approach and explores new sellers at a rate proportional to the ratio of unexplored sellers to all the sellers who have sent price bids. For phases I to III, we considered just one buyer attitude, as the focus was on the evaluation of the trust and reputation model. As risk taking buyer is more aggressive in trying new sellers, we used risk taking buyers in phases I to III.

In Phase IV, buyers with risk taking and conservative approaches are analyzed. In Phase IV, the market was populated with various sellers (honest, dishonest, varying quality, and varying price) and different types of friends (honest, lying, sharing similar opinions, and having different opinions). We compared the performances of risk taking and conservative buyers when they are acting alone and when they have friends. Metric used was: percentage of purchases from honest, high quality sellers at different levels of purchase, and as the number of sellers in the market was varied. We wanted to see how attitude affects performance of buyers when they are acting alone and when they have friends.

In the last phase, Phase V, we compared performances of all buyers (risk taking, conservative, acting alone, and with friends) at different levels of purchases, and as the number of sellers was varied. We considered a market populated with various sellers (honest, dishonest, varying quality, and varying price) and different types of friends (honest, lying, sharing similar opinions, and having different opinions). We compared all buyers on the following metric: percentage of purchases from honest, high quality sellers at different levels of purchase, and as the number of sellers in the market was varied. We wanted to see if any particular combination of buyer attitude and seller modeling yields the best performance.

For our experiments we developed a simulation of an electronic market. Figure 4.1 provides a simplified overview of the market's architecture. The simulation consists of a Matchmaker [31], buyer agents, and seller agents, which communicate via KQML messages [17].

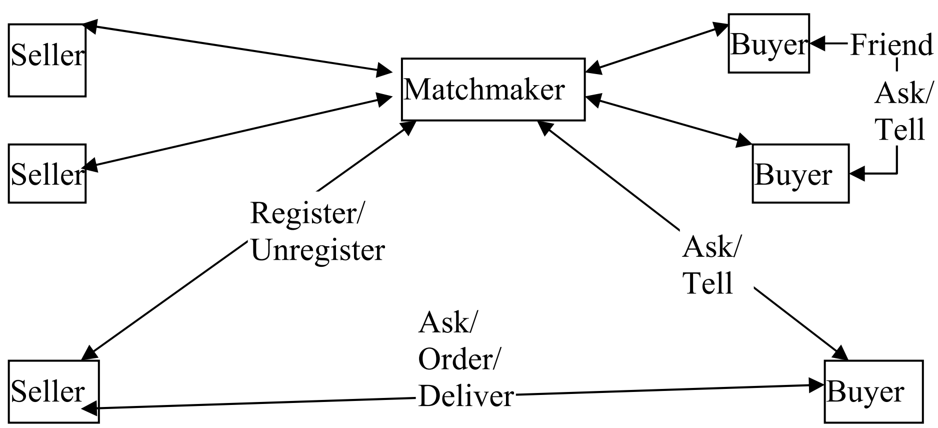

Figure 4.1: Market simulation comprising of Matchmaker, buyer, and seller agents.

Seller agents join the market by registering with the Matchmaker, informing it of the products they are offering and their prices. They also notify the Matchmaker before exiting the market by un-registering. Thus, the Matchmaker has a current list of all sellers active in the market. Buyers send *ask* messages to the Matchmaker to get information about sellers for a certain product. Matchmaker responds with a *tell* message containing information about all the sellers selling that product. Buyers then send *ask* messages to sellers requesting price quotes for the product. Upon receiving an order a seller responds via a *deliver* message which contains a value representing the quality of the product. A buyer then uses the quality and the price charged by the seller to calculate a trust rating for the seller for that interaction. Each buyer maintains a private record for all seller agents from whom it has purchased products in the past and utilizes this information to compute average direct trust rating as discussed in section 3.1 in chapter 3.

Buyers may belong to a community comprising of buyers. Members of the community are referred to as *friends*. While evaluating sellers for purchase decisions, a buyer sends *ask* messages to friends requesting for seller information. Friends respond via *tell* messages containing information about sellers. For our experiment in Phase I, buyers do not any have friends. For experiments in phases II to V, buyers belong to a community and have friends.

All the messages exchanged between agents have the following format:

action from to sender-type ontology content

Action specifies the nature of message and it can be *ask*, *tell*, *deliver*, *order*, *register,* and *unregister*. From specifies message sender's id. To specifies the recipient of message. Sender-type specifies whether sender is a seller, buyer or matchmaker. Ontology specifies how the content field of message is to be interpreted. Content gives the actual content of message.

For example, when seller *S1* enters the market and wants buyers to be aware of it as a seller for products *P1* and *P2*, it sends a *register* message to Matchmaker *M1*.

register S1 M1 seller default P1 P2

When seller *S1* want to leave the market it sends an *unregister* message to *M1*.

unregister S1 M1 seller default P1 P2

When buyer *B1* wants to get a list of sellers selling product *P1*, it sends the following *ask* message to *M1*.

ask B1 M1 buyer seller-list P1

Matchmaker responds by sending a *tell* message with list of sellers *(S1, S3, S7)* selling product *P1.*

tell M1 B1 matchmaker seller-list S1 S3 S7

The following message is an example of buyer *B1* requesting for a price quote from seller *S1* for product *P1*.

ask B1 S1 buyer price P1

Seller *S1* responds by sending a *tell* message which contains price.

tell S1 B1 seller price P1 45

In the following message, buyer *B1* places an order for product *P1* with seller *S1*.

order B1 S1 buyer default P1

In the following message, seller *S1* delivers product *P1* of quality 47 to buyer *B1*

deliver S1 B1 seller default 47

Buyer B1 requests for seller information on sellers S1, S3 and S7 from buyer B2 by sending the following ask message:

ask B1 B2 buyer seller-reputation S1 S3 S7

Buyer *B2* responds with the following tell message:

tell B2 B1 buyer seller-reputation S1 recommend 0.8 34 S3 neutral 0 0 S7 notrecommend 0.1 5

4.1 Phase I

Our aim in Phase I was to compare the performance of our model of rating sellers based on direct interactions with other seller evaluating strategies, based only on direct interactions that have been proposed for this kind of market.

In our experiment, we compared the performances of four buyers. All four buyers make decisions regarding sellers based on their own experiences and do not have any friend buyers.

1. *Risk Buyer*: - This buyer uses our direct trust rating model for evaluating sellers as described in section 3.1 in chapter 3 and has risk taking attitude.

2. *Tran Buyer*: - This buyer uses buying strategy as described in Tran and Cohen [49, 50] and summarized in section 2.4.1 in chapter 2.

3. *RL Buyer*: - This buyer uses a reinforcement learning strategy as described for 0-level buying agent in Vidal and Durfee [53] and summarized in 2.3.3 in chapter 2.

4. *Random Buyer*: - This buyer chooses a seller randomly.

For all buyers, we set buyers' product valuation function as *3 * quality – price* to reflect that high quality is buyers' first preference. A buyer's product valuation function reflects the gain a buyer makes from having purchased a product from a seller. The range of quality *q* sold across sellers was randomly set from 10 to 50 and varied in units of 1. We set Risk Buyer's and Tran Buyer's minimum quality expectation to 40 to indicate that both buyers desire high quality product. Hence acceptable quality for Risk Buyer and Tran Buyer is from 40 to 50 and non acceptable quality is from 10 to 39. They use this to rate sellers, and sellers with high ratings are identified as potential sellers. A potential seller who maximizes product valuation is then chosen as seller to purchase from. RL Buyer and Random Buyer do not have a minimum quality expectation. RL Buyer chooses a seller who maximizes its product valuation, and Random Buyer chooses a seller randomly.

Price, *pr,* of a product for honest sellers was set to $q \pm 10\%q$. Like Tran [49], we make the assumption that it costs more to produce high quality goods. We also make a reasonable assumption that seller may offer a discount to attract buyers in the market or raise its price slightly to increase its profits. Hence price charged by an honest seller was set from 90% to 110% of product quality. Dishonest sellers charge higher prices. To compare various buyers, we populated the market with sellers offering different qualities at various prices. Based on quality delivered, price charged, and nature of sellers we considered the following six categories of sellers:

1. *Honest Acceptable (HA):* - Each seller offers quality in the buyer's acceptable range [40, 50]. Their price is from 90% to110% of the quality they are selling.

2. *Honest Not Acceptable (HNA):* - Each seller offers a quality in the buyer's unacceptable range [10, 39]. Their price is from 90% to110% of the quality they are selling.

3. *Overpriced Acceptable (OPA):*- Each seller offers quality in the buyer's acceptable range [40, 50]. Their price is from 111% to 200% of the quality they are selling.

4. *Overpriced Not Acceptable (OPNA):* - Each seller offers quality in the buyer's unacceptable range [10, 39]. Their price is from 111% to 200% of the quality they are selling.

5. *Inconsistent:* - Each seller offers quality in the range [10, 50]. Their price is from 90% to 110% of the quality they are selling.

6. *Dishonest:* - This category of sellers in their first sale to a buyer offer acceptable quality q [40, 50] charging a price $pr = q \pm 10\%q$. In their subsequent sales to that buyer they reduce the quality q to be in the range [10, 25]. However their price still remains high, by setting it to $pr = q1 \pm 10\%q1$ where $q1$ is in the range [40, 50].

Table 4.1 lists various seller categories along with their price and quality characteristics.

Table 4.1: Seller quality and price configuration

Seller Category	Quality q	Price
Honest Acceptable	40-50	(0.9-1.1) * q
Honest Not Acceptable	10-39	(0.9-1.1) * q
Overpriced Acceptable	40-50	(1.11-2.00) * q
Overpriced Not Acceptable	10-39	(1.11-2.00) * q
Inconsistent	10-50	(0.9-1.1) * q
Dishonest First Sale Subsequent Sales	 q_{first} = 40-50 10-25	 (0.9-1.1) * q_{first} (0.9-1.1) * q_{first}

Parameters for Risk Buyer: As mentioned previously Risk Buyers' minimum quality expectation was set to 40. Risk Buyer's desired expected quality q_{exp} was set to *50*. This value reflects highest quality that buyer desires and should be set to a value greater than or equal to the buyer's minimum quality expectation. As the maximum quality delivered by any seller is 50, maximum price p_{max} quoted by an honest seller would be 55(50+10%*50). As the minimum quality delivered by any seller is 10, the minimum price p_{min} quoted would be 9 (10-10%10). With an even distribution, the average price p_{avg} would be 32 (pmax+pmin/2). The threshold values θ for a seller to be considered reputable and ω for a seller to be considered disreputable values can be computed as follows: Risk Buyer is expecting at least a quality of 40. In the worst case, it can get this at the highest price that can be charged by an honest seller which would be 44(40+10%*40). From equation 1(a) the trust rating for that seller would be

$$\frac{40}{50} - \left(\frac{44 - 32}{55} \right) = 0.581$$

So we set $\theta = 0.58$. For new sellers trust rating is set to 0. These buyers should not come under the category of disreputable sellers. So the threshold value, ω, for a seller to be considered unacceptable should be less than 0 and we set it as -0.1.

Parameters for Tran Buyer: We set this buyer's minimum quality expectation to 40, same as Risk Buyer's minimum quality expectation. The threshold for a seller to be considered reputable is set to 0.5 and for a seller to be considered disreputable is set to -0.9 as described in their work.

RL Buyer and Random Buyer do not have specific parameters to be set. RL Buyer chooses a seller who maximizes its product valuation, and Random Buyer chooses a seller randomly.

We wanted the market to have at least one seller from each seller category shown in table 4.1. To keep seller distribution even, we wanted the market to have equal number of sellers from each category. This means the market should have 6 sellers at the minimum. We populated the market with two sellers per category, a total of 12 sellers. In our experiment each buyer conducted 500 transactions. As there are 12 sellers, this should give buyers sufficient time to identify desirable sellers to purchase from and then consistently buy from these sellers. In each transaction, each buyer purchases product *p* by querying the seller list from the Matchmaker, obtaining price quotes from different sellers and utilizing its buying strategy to choose a seller. Data shown is based on the average of 100 runs of the experiment.

In our experiment we compared the performances of various buyers on the following parameters:

- How long it took them to learn to identify high quality, low priced sellers. We want buyers to identify high quality sellers offering low prices as soon as possible. If a buyer is able to identify high quality sellers quickly, then the buyer's strategy can be used when making infrequent purchases.

- Average gain as the number of purchases of product p is increased. If the average gain is consistently high, it means that the buyer is interacting with high quality sellers offering low prices most often. If the average gain is high earlier on, it implies that buyer has identified high quality low price sellers quickly.

Figures 4.2 - 4.5 show gain versus transactions for each type of buyer.

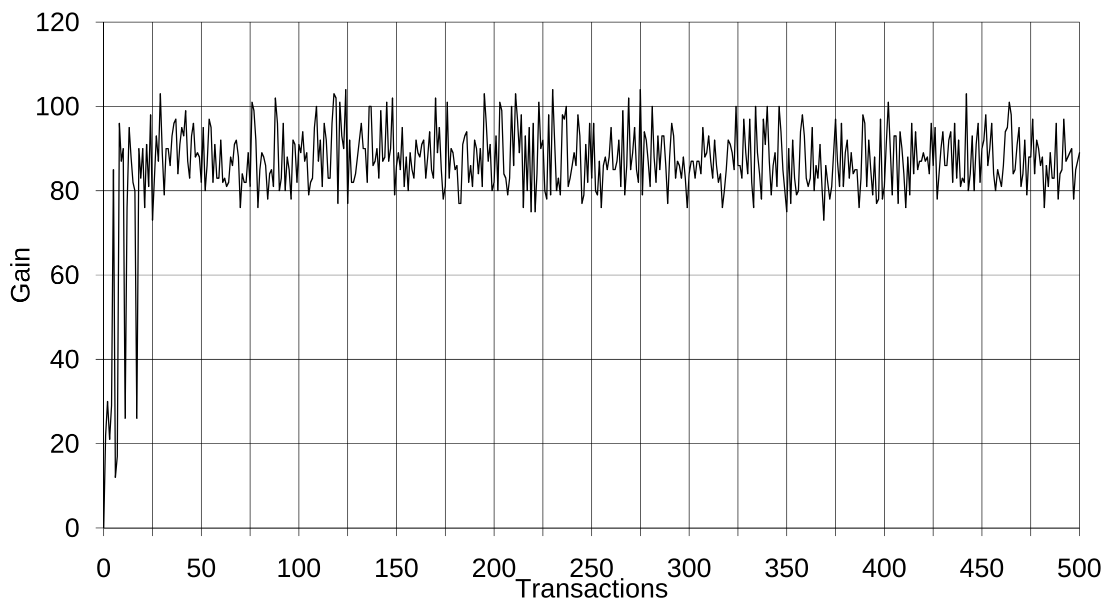

Figure 4.2: Gain versus transactions for Risk Buyer.

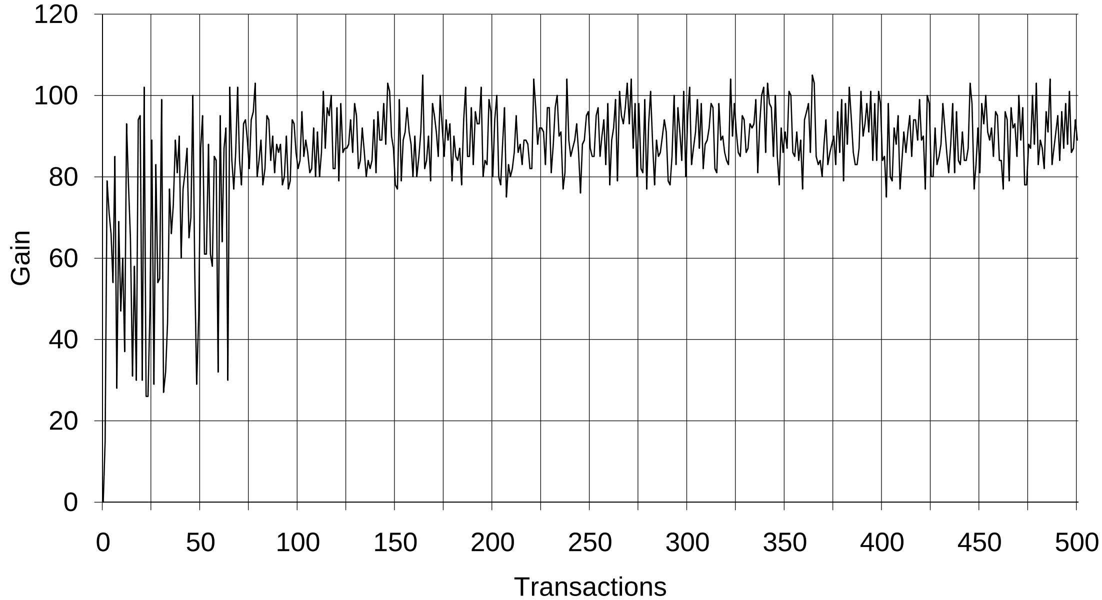

Figure 4.3: Gain versus transactions for Tran Buyer.

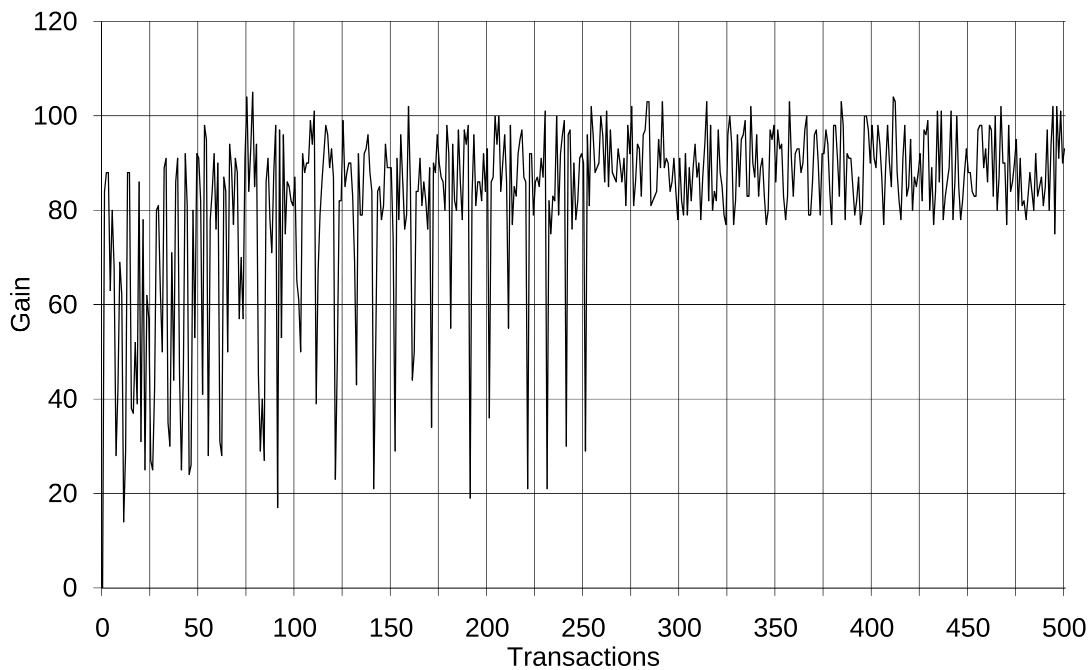

Figure 4.4: Gain versus transactions for RL Buyer.

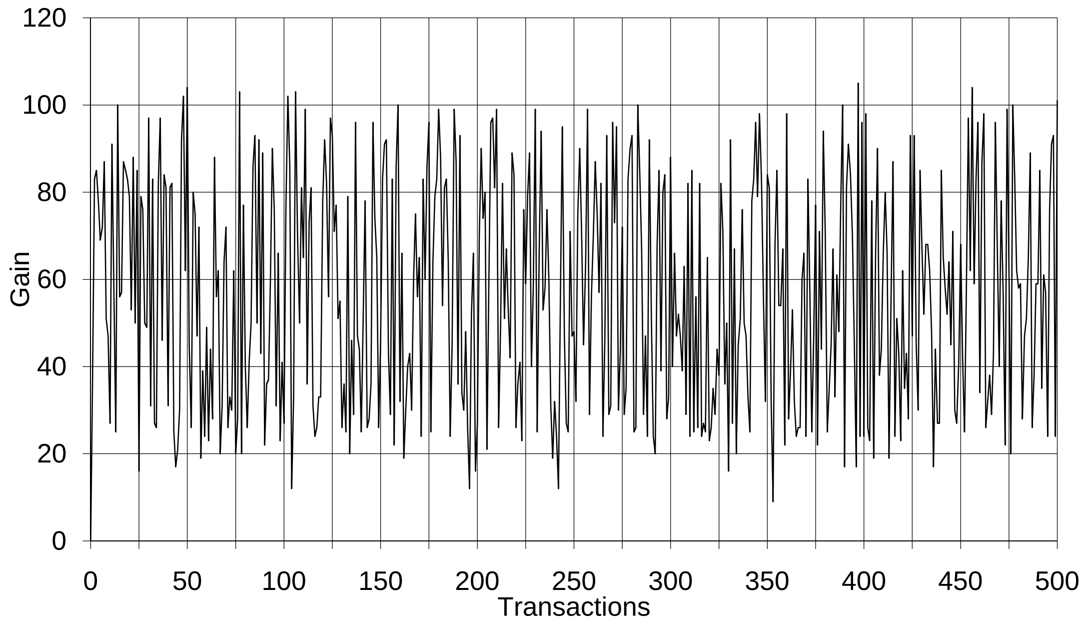

Figure 4.5: Gain versus transactions for Random Buyer.

Table 4.2 shows the number of purchases made by a buyer from each seller type.

Table 4.2: Buyer-seller interactions

	HA	HNA	OPA	OPNA	INC	DIS
Risk Buyer	488	2	2	2	2	4
Tran Buyer	451	7	23	5	8	6
RL Buyer	420	16	15	13	17	16
Random Buyer	86	88	82	83	69	92

Acceptable quality sellers can offer qualities anywhere from 40 to 50. The lowest gain from purchasing from an honest seller offering at the lowest end of good quality range and charging its highest price is 76 (3*40 – 44). When the gain from purchasing from a seller is 76 and above, it means the buyer is purchasing from a high quality low priced seller. From figures 4.2-4.5, it can be seen that Risk Buyer, Tran Buyer, and RL Buyer learn at different rates to identify high quality, low priced sellers. After having learnt, they consistently interact with high quality, low priced sellers. This is confirmed by the fact that highest number of purchases are made from honest acceptable sellers as shown in table 4.2. Random Buyer never learns and that is to be expected as it is choosing sellers randomly. Risk Buyer learns to identify high quality low priced sellers very quickly in about 15 transactions or purchases, Tran Buyer takes about 60 transactions to learn, and RL Buyer learns in about 250 transactions. If the buyers were to purchase the product infrequently then our model would work better than RL Buyer's or Tran Buyer's strategy as it requires the least number of transactions to learn.

Figure 4.6 shows average gain versus the number of purchases for different buyers.

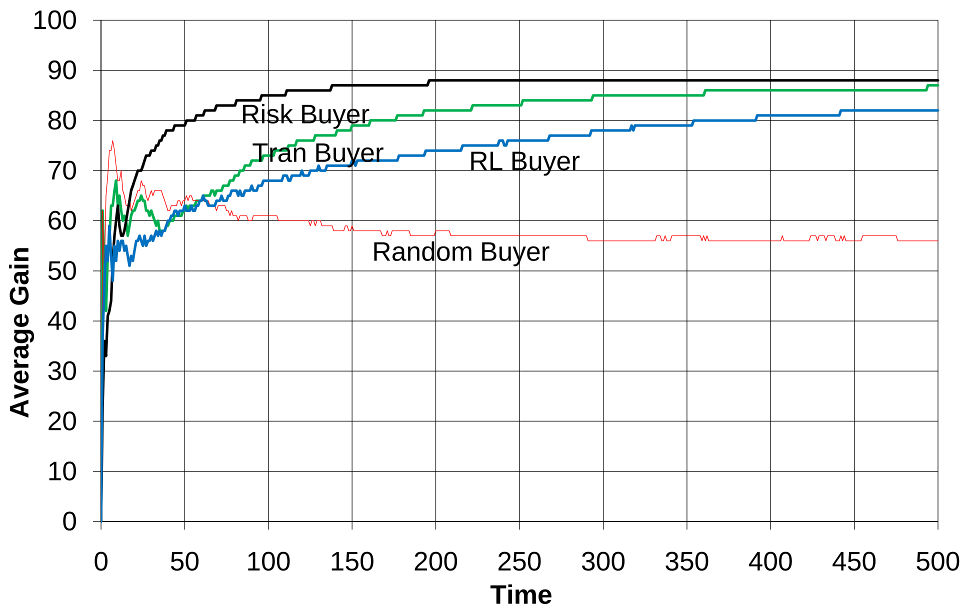

Figure 4.6: Average gain versus number of purchases for different buyers.

In the beginning, average gains are fluctuating as buyers employing a non-random strategy are learning and Random Buyer is choosing sellers randomly. Risk Buyer is the quickest to learn and its average gain raises sharply earlier on compared to the other two learning agents. As RL Buyer takes a long time to learn, its average gain at the end is still lower than the Risk Buyer's or Tran Buyer's. Since Random Buyer purchases randomly from various types of sellers, its average is consistently the lowest. In the first half of figure 4.6 it can be seen that when purchases are fewer, average gain for Risk Buyer, once its learning phase is completed, is higher than other buyers. So, if buyers were to purchase a product infrequently, then our model works better than RL Buyer's or Tran Buyer's strategies. As the number of purchases increases, Risk Buyer still has the highest average gain with Tran Buyer's average gain coming very close to it at very high number of purchases.

4.2 Experimental setup for Phases II, III, IV, and V

In Phase II, our aim was to analyze the first method of identifying friends presented in this dissertation. In Phase III, we wanted to compare the first and second method of identifying friends. In Phase IV, we wanted to study risk and conservative attitudes of buyers, and in Phase V, we wanted to see if any particular combination of buyer attitude and modeling of sellers yields the best performance. In the electronic market simulation shown in figure 4.1, we populated the market with:

- sellers of different nature (honest, dishonest, and inconsistent) selling different qualities (high, medium, and low) at different prices,

- buyers with different attitudes (risk taking and conservative) either acting alone , or with friends utilizing one of the methods of identifying friends,
- honest friend buyers who either share same opinions, or have slightly different opinions, or very different opinions, and,
- lying friend buyers.

We conducted two experiments and analyzed the performance of various buyers in different phases. The following sections describe the sellers, buyers, friend buyers and the experiments in detail.

4.2.1 Sellers

Similar to our settings in Phase I, for our experiments in Phase II, sellers sell a single product and its quality q ranges from 10 to 50 and varies in units of one. We categorized quality into three ranges: high (40-50), medium (25-39) and low (10-24). In Phase I, we had categorized quality in acceptable and unacceptable ranges. Here we categorized them into three ranges because we wanted to simulate buyers with different opinions regarding sellers. Similar to our settings in Phase I, for our experiments in Phase II, the price charged for a product by a seller depends on the nature of that seller; this can be honest, dishonest, inconsistent, or overcharging. Like Tran [49, 50], we make the assumption that it costs more to produce high quality goods and that the seller may offer a discount to attract buyers in the market or raise its price slightly to increase its profits. Hence the price of the product is set in the range of 90% to110% of the quality value for an honest seller. Overpriced sellers charge 111% to 200% of the quality value. Each honest seller offers the product in any one quality category. Dishonest sellers offer high quality products and charge an honest seller's price in their first transaction with a buyer; subsequently, they reduce the quality by 50%, but continue to charge the same price. Inconsistent sellers sell high, medium and low quality products randomly and charge honestly. For our experiments, we considered sellers belonging to one of the following eight categories created by different combinations of seller nature and quality ranges with price and quality values as shown in Table 4.3.

Table 4.3: Seller quality and price configuration

Seller Category (Nature Quality)	Quality q	Price
Honest High	40-50	(0.9-1.1) * q
Honest Medium	25-39	(0.9-1.1) * q
Honest Low	10-24	(0.9-1.1) * q
Overpriced High	40-50	(1.11-2.00) * q
Overpriced Medium	25-39	(1.11-2.00) * q
Overpriced Low	10-24	(1.11-2.00) * q
Inconsistent	10-50	(0.9-1.1) * q
Dishonest First Sale Subsequent Sales	 q_{first} = 40-50 20-25	 (0.9-1.1) * q_{first} (0.9-1.1) * q_{first}

Similar to our settings in Phase I, for our experiments in Phase II we set the product valuation function to be *3*quality – price*, which implies that purchasing highest quality product is the buyer's first priority, and the buyer would like to purchase the highest quality product from a seller who is selling it at the lowest price, i.e., the category of sellers belonging to "Honest High".

4.2.2 Buyers

We compared the performances of the following six buyers:

1. *Risk Buyer:* This buyer makes decisions regarding purchases based on its own experience with sellers and uses our direct trust rating model as described in section 3.1 in chapter 3 for evaluating sellers. Its attitude towards new buyers is "Risk Taking" which means that new or unexplored buyers are always included in the list of potential sellers. Similar to our settings in Phase I, this buyer expects a minimum quality of 40 and its desired expected quality is 50. The buyer's reputation rating scale for sellers is [-1, 1]. The reputation threshold values θ and ω for a seller to be considered reputable and disreputable respectively were calculated as described in the Phase I and are set to 0.6 and -0.1. This buyer's product valuation is maximized by purchasing from sellers belonging to the "Honest High" category and, consequently, sellers from this category get the highest rating by this buyer.

2. *RiskWithFr Buyer:* A buyer with friends using our direct trust rating model for evaluating sellers based on direct interactions, a "Risk Taking" attitude towards new buyers, and using the first method of identifying trustworthy friends as described in section 3.2.1 in chapter 3. This buyer uses friends' opinions as well and friends' ratings of sellers, when considering various sellers for purchase decisions. The quality expectations from and reputation thresholds for a seller are the same as for a Risk Buyer and hence this buyer's

product valuation is also maximized by purchasing from the "Honest High" category of sellers.

3. *OOnlyRiskWithFr Buyer:* A buyer with friends, using our direct trust rating model for evaluating sellers (section 3.1) based on direct interactions, "Risk Taking" attitude towards new buyers, and using the second method of identifying trustworthy friends as described in section 3.2.2 in chapter 3. This buyer considers only friends' opinions regarding sellers when making decisions about sellers. The quality expectations from and reputation thresholds for this seller are the same as for a RiskWithFr Buyer and hence this buyer's product valuation is also maximized by purchasing from the "Honest High" category of sellers.

4. *Cns Buyer*: This buyer makes decisions regarding purchases based on its own experience with sellers and uses our direct trust rating model for evaluating sellers. Its attitude towards new buyers is "Conservative". This buyer purchases from new sellers during its exploration of the market. The exploration rate is proportional to the ratio of new sellers to all the sellers in the market. The detailed policy of a buyer with conservative attitude regarding new sellers is described section 3.5 of chapter 3. The desired expected quality, acceptable quality limits and threshold values for θ and ω are the same as for Risk Buyer.

5. *CnsWithFr Buyer:* A buyer with friends using our direct trust model for evaluating sellers based on direct interactions, having "Conservative" attitude towards new buyers and using the first method of identifying trustworthy friends as described in section 3.2.1 in chapter 3. The quality expectations from, and reputation thresholds for a seller are the same as for a Cns Buyer and hence this buyer's product valuation is also maximized by purchasing from the "Honest High" category of sellers.

6. *OOnlyCnsWithFr Buyer:* A buyer with friends using our direct trust rating model for evaluating sellers based on direct interactions (section 3.1), having "Conservative" attitude towards new buyers and using the second method of identifying trustworthy friends as described in section 3.2.2 in chapter 3. The quality expectations from, and reputation thresholds for a seller are the same as for a RiskWithFr Buyer and hence this buyer's product valuation is also maximized by purchasing from the "Honest High" category sellers.

Table 4.4 lists various buyers along with their distinguishing features whose performances were analyzed in different phases.

Table 4.4: Different buyers' characteristics

Buyer	Attitude	Friends	Friend Model
Risk Buyer	Risk taking	None	None
RiskWithFr Buyer	Risk taking	Yes	Utilize friend's opinion and numerical rating of a seller
OOnlyRiskWithFr Buyer	Risk taking	Yes	Utilize friend's opinion of a seller
Cns Buyer	Conservative	None	None
CnsWithFr Buyer	Conservative	Yes	Utilize friend's opinion and numerical rating of a seller
OOnlyCnsWithFr Buyer	Conservative	Yes	Utilize friend's opinion of a seller

In Phase I, we showed that a "Risk Buyer" has a superior performance than agents employing Tran's [49] model (Tran Buyer), or Vidal and Durfee's [53] model (RL Buyer) for frequent as well as for infrequent purchases. Hence we did not include Tran Buyer or RL Buyer in our experimental phases II-IV

4.2.3 Friend Buyers

We considered three categories of friends:

4.2.3.1 Friends with similar opinions

The first category is a group of friends who share similar opinions, but some of them have different seller rating systems than RiskWithFr Buyer. All buyers belonging to this group have the same quality expectations as RiskWithFr Buyer; give high ratings to sellers belonging to "Honest High" category and respond honestly when requested for information. We considered four friends in this category:

1. *Risk Friend*: This friend uses our direct trust rating model for evaluating sellers for direct interaction. This friend is identical to RiskWithFr Buyer in its method of rating sellers based on direct interactions.

2. *CD-0.5 Friend*: This friend uses a scale that is additively different from RiskWithFr Buyer's scale by 0.5. This buyer's reputation rating scale for rating sellers is [-0.5, 1.5]. The CD in the buyer's name stands for "constant difference" as the difference in the reputation rating of a seller between this buyer and RiskWithFr Buyer is always 0.5.

3. *CM-10 Friend:* This friend has a rating scale which is ten times that of the RiskWithFr Buyer. This buyer's reputation rating scale for rating sellers is [-10, 10]. The CM in the buyer's name stands for "constant multiple" as the reputation rating of a seller by this buyer is a constant multiple (10 times) of the reputation rating of the same seller by RiskWithFr Buyer.

4. *N-01 Friend:* This friend's seller rating scale is [0, 1]. The N in the buyer's name stands for "normalized". We implemented this buyer by taking Risk Friend buyer, whose reputation rating scale is from -1 to 1, and normalized the reputation values to fit on a [0, 1] scale.

4.2.3.2 Friends with slightly different opinions

Second category of friends has slightly different opinions and responds honestly when asked for information. We considered two buyers in this category.

5. *SDO Friend:* This buyer's rating system for sellers is similar to RiskWithFr Buyer's rating system for direct interactions; however, the quality expectations are slightly different. This friend is satisfied with medium quality products, does not care for higher quality, and gives high rating to sellers belonging to the "Honest Medium" category. Sellers belonging to "Honest High" category are given slightly lower ratings than sellers belonging to "Honest Medium" category, as this quality product is slightly more expensive than medium quality product. This buyer's reputation rating scale for rating sellers is [-1, 1]. SDO in the buyer's name stands for "Slightly Different Opinion".

6. *Tran Friend:* This friend uses Tran's model for evaluating sellers [49]. This buyer gives high ratings to sellers belonging to "Honest High' category as its quality expectations are same as RiskWithFr Buyer. This buyer's opinions are slightly different as it classifies sellers as good or bad; where as RiskWithFr Buyer classifies sellers as good, neutral, and bad. This buyer's reputation rating scale for rating sellers is [-1, 1].

4.2.3.3 Friends with very different opinions

This category of friends has very different opinions, but is honest when responding to information requests. We considered one friend in this category.

7. *VDO Friend:* This friend is satisfied with low quality sellers and does not care for quality to be greater than 24, which is the upper bound of the low quality range. It gives sellers belonging to "Honest Low" category high ratings. Sellers belonging to "Honest Medium" category are given slightly lower ratings than sellers belonging to "Honest Low" category, as they are considered slightly expensive, and sellers belonging to "Honest High" category are given low ratings as they charge higher prices. Sellers belonging to "Honest Low" and sometimes "Honest Medium" quality categories are considered as good sellers. Sellers belonging to "Honest High" quality category are not considered as good sellers, and this friend does not recommend this category of sellers. VDO in the buyer's name stands for "Very Different Opinion". This buyer's reputation rating scale for rating sellers is [-1, 1].

4.2.3.4 Lying friends

These types of friends are not honest when responding to information request. We considered one friend in this category.

8. *Erratic Friend:* This buyer's rating system for sellers is similar to RiskWithFr Buyer's rating system for direct interactions. Its reputation rating scale is [-1,

1]. Replies from this friend to an information request are erratic. The reputation of the sellers is provided truthfully and incorrectly on a random basis.

Table 4.5 summarizes the various types of friends we considered.

Table 4.5: Friends' characteristics

Friend Buyer	Direct Interaction Model	Opinion Similarity	Honesty	Rating Scale
Risk Friend	Our direct trust rating model	Very Similar	Honest	[-1,1]
CD-0.5 Friend	Our direct trust rating model	Very Similar	Honest	[-0.5, 0.5]
CM-10 Friend	Our direct trust rating model	Very Similar	Honest	[-10, 10]
N-01 Friend	Our direct trust rating model	Very Similar	Honest	[0, 1]
SDO Friend	Our direct trust rating model	Slightly different	Honest	[-1, 1]
Tran Friend	Tran's Model	Slightly different	Honest	[-1, 1]
VDO Friend	Our direct trust rating model	Very different	Honest	[-1, 1]
Erratic Friend	Our direct trust rating model	Very Similar	Lying	[-1, 1]

4.2.4 Parameters

The threshold value ψ for a friend to be considered reputable was set to 0.6, the exact value as θ, the reputation threshold value for a seller to be considered reputable. We set the value of α and β to be 0.2 and -0.6 respectively (see equations (5, 6)). When $\alpha > 0.5$, a friend's reputation value can cross the threshold value within one transaction, which is too soon to determine a friend's trustworthiness, reliability and opinions. So α should be between [0.1, 0.5], and for a particular value of α, $\beta > \alpha$. The reason being, after becoming a trustworthy friend, if the reputation of the friend falls below the threshold value ψ, it should drop sufficiently that its reputation value does not cross the threshold value ψ within one transaction. Initially we set α = 0.2 and β = -0.6 and later conducted experiments to determine the optimum range of values for α and β. With α = 0.2, initially it will take at least 6 transactions for a friends' reputation to exceed the threshold value and hence we set m in equations (8 and 9) to be 5, and the standard deviation threshold μ to be 0.15. We set the value of x and y to be 0.5 each, which means the buyer, will give 50% weight to its own experiences and 50% weight to the opinions of its friends (see equation (13)).

4.2.5 Experiments

We conducted two experiments. In experiment 1, we populated the market with 32 sellers, with four sellers from each category in table 4.3. We wanted to have the same number of sellers in each category. Since there are eight categories, we had to populate the market with sellers in multiples of 8. For our first experiment, we chose 32 sellers and in experiment 2 we varied the number of sellers. We populated the market with 14 buyers comprising of 6 buyers, one from each category in table 4.4 and one of each type of 8 friends summarized in table 4.5. We studied the performances of six buyers chosen from each category in table 4.4. To ensure that the model works correctly and consistently, we conducted 100 trials and report on the average of 100 trials. In each trial, the buyers made 500 purchases or transactions with sellers. 500 transactions would give buyers ample time to learn and purchase from high quality, low priced sellers constantly. In each transaction, each buyer obtains a registered list of sellers selling this product from the Matchmaker and sends a message to each seller in the list asking them to submit their bids for the product. Sellers who are interested in getting the contract submit a bid which includes the price. Buyers with friends then request information about sellers from each buyer friend. All buyers wait for a certain amount of time for responses and then evaluate the bids received to choose a seller to purchase from. We had the friends make their purchases first so that they could provide useful information to buyers with friends.

In experiment 2, the buyer configuration was the same as in experiment 1. We varied the number of sellers in the market and conducted 100 trials for each seller count. Buyers made 500 transactions in each trial. In each transaction, the buyers made purchases as described in experiment 1. As the number of sellers was increased, we kept the ratio of each type of seller to the total number of sellers constant. For example, when the market was populated with 8 sellers, there was one seller from each category in Table 4.3 and when the seller count was increased to 32, there were 4 sellers from each category.

Using the data collected from experiment 1 and 2, the performances of various buyers was analyzed in different phases. In Phase II, the first method of identifying trustworthy friends as described in section 3.2.1 was analyzed by studying the performance of RiskWithFr Buyer. RiskWithFr Buyer considers a trusted friend's opinion and rating of a seller in choosing a seller to purchase from. In Phase III, the first and the second methods of identifying trusted friends were analyzed by comparing the performances of RiskWithFr Buyer and OOnlyRiskWithFr Buyer.

OOnlyRiskWithFr Buyer considers only trusted friends' opinions of a seller. In Phase IV, risk taking and conservative attitudes of buyers were compared and in Phase V all the six buyers in table 4.4 were compared.

4.3 Phase II

In Phase II, the first method of identifying trustworthy friends as described in section 3.2.1 in chapter 3 was analyzed for its effectiveness 1) in successfully identifying trustworthy friends, and 2) utilizing the information regarding sellers (friend's opinion regarding the seller and the numerical rating for that seller) provided by these trustworthy friends along with its own information in making purchase decisions. We compared the gains of a buyer acting alone (Risk Buyer) and a buyer with friends (RiskWithFr Buyer) to determine if the first method presented in section 3.2.1 helps in improving the performance of buyer with friends as compared to a buyer making decisions based on its own information.

The experimental setup and the two experiments have been described in the previous section (4.2). From the data collected from experiments 1 and 2, in this phase, we analyzed the performances of Risk Buyer and RiskWithFr Buyer. Risk Buyer is a buyer acting alone and making decisions regarding sellers based on its own information. It uses the direct trust rating model as described in section 3.1 to model sellers. RiskWithFr Buyer uses our direct trust rating model for direct interactions, has friends, and uses the first method of identifying trustworthy friends. In the first method, friends' opinions and the ratings of sellers are used to identify trustworthy friends (who are honest and share similar opinions). Next, a buyer adjusts sellers' ratings by trustworthy friends to account for the differences in the buyer's and trustworthy friends' seller rating systems and then combines them with its ratings of sellers to evaluate sellers for purchase decisions.

Our first goal was to study the effectiveness of the first method of identifying trustworthy friends. From the data collected for RiskWithFr Buyer in experiment 1: 1) we looked at how many times different types of friends were identified as trustworthy in 100 trials, 2) in each trial, how soon was each type of friend identified as trustworthy, and 3) in each trial, the number of times a friend's opinion was utilized out of 500 transactions. We used these parameters for comparison because a friend is identified as trustworthy if the opinions of RiskWithFr Buyer and the friend regarding a seller match, and the standard deviation (of the differences between the two buyers' ratings for a seller) is below a threshold value. Since friends are being monitored continuously, how often a friend's opinion is used will depend on the type of friend. We expected friends with similar opinions (Risk Friend, CD-0.5 Friend, CM-10 Friend, and N-01 Friend) to be identified the earliest and their opinions to be utilized all the time once they are identified. We expected friends with slightly different opinions (SDO Friend, and Tran Friend) also to be identified as trustworthy friends a little later and their opinions to be utilized. To ensure that the model works correctly and consistently, we conducted 100 trials and expected similar and slightly different opinion friends to be identified as trustworthy in all of the trials. Ideally, VDO Friend and Erratic Friend should not be identified as trustworthy in any of the trials.

The results from experiment 1 are shown in table 4.6.

Table 4.6: Performance of buyer using first method of identifying trustworthy friends

Friend	Average transaction when a friend is first identified as a trustworthy friend in each trial[1]	Average percentage of 500 transactions a friend's opinion is used in each trial.	Percentage of trials identified as a trustworthy friend in 100 trials
Friends with similar opinions			
Risk Friend	13	96%	100%
CD-0.5 Friend	8	96%	100%
CM-10 Friend	39	76%	100%
N-01 Friend	11	94%	100%
Friends with slightly different opinions			
SDO Friend	28	84%	100%
Tran Friend	119	72%	100%
Friends with very different opinions			
VDO Friend	38	3%	14%
Lying Friend			
Erratic Friend	20	0.4%	3%

[1] A friend is identified as trustworthy only if the opinions of RiskWithFr Buyer and the friend regarding a seller match and the standard deviation of the differences between the buyers' ratings for sellers is below a threshold value.

From Table 4.6 we see that friends with similar opinions were identified as trustworthy in all the trials, they were among the earliest to be identified as trustworthy, and their opinions were utilized 76% to 96% of 500 transactions.

SDO Friend was also identified as trustworthy in all the trials and its opinions were utilized in 84% of 500 transactions in each trial. The second "slightly different" friend, Tran Friend, was also identified as a trustworthy friend in all the trials, and its opinions were utilized in 72% of 500 transactions. It was identified as trustworthy significantly later than other trustworthy friends. We think this is because Tran Friend classifies sellers as good and bad (whereas our model classifies sellers as recommend, not recommend, and no opinion) and hence it takes a longer time for RiskWithFr Buyer to identify Tran Friend as trustworthy.

VDO Friend was correctly identified as untrustworthy in 86% of the trials. It was incorrectly identified as trustworthy in 14% of the trials; however, its opinions were utilized in only 3% of 500 transactions in each trial it was identified as trustworthy. After erroneously identifying VDO Friend as trustworthy early on, RiskWithFr Buyer is able to quickly realize that this friend should not be considered trustworthy and removes it from its trustworthy friends list. This is reflected in the fact that the friend's opinions were utilized in only 3% of 500 transactions.

Erratic Friend, who provides information about sellers randomly, was correctly identified as not trustworthy in 97% of 100 trials. In 3% of the trials, even though the model incorrectly identified this buyer as a trustworthy friend, RiskWithFr Buyer is able to realize its mistake which is indicated by the fact that this friend's opinion was utilized in just 0.4% of all 500 transactions.

Our next goal was to determine whether a buyer with friends, utilizing the first method of identifying trustworthy friends, has higher performance than a buyer acting alone. From the data collected in experiment 1, we compared the average percentage of purchases made from honest, high quality sellers by Risk Buyer and RiskWithFr Buyer at different levels of purchases.

Figure 4.7 shows the average percentage of purchases made from honest, high quality sellers by Risk Buyer who is acting alone and RiskWithFr Buyer who is utilizing the opinions of friends along with its own experiences in making decisions at different levels of purchases, as the buyers made 500 purchases in total.

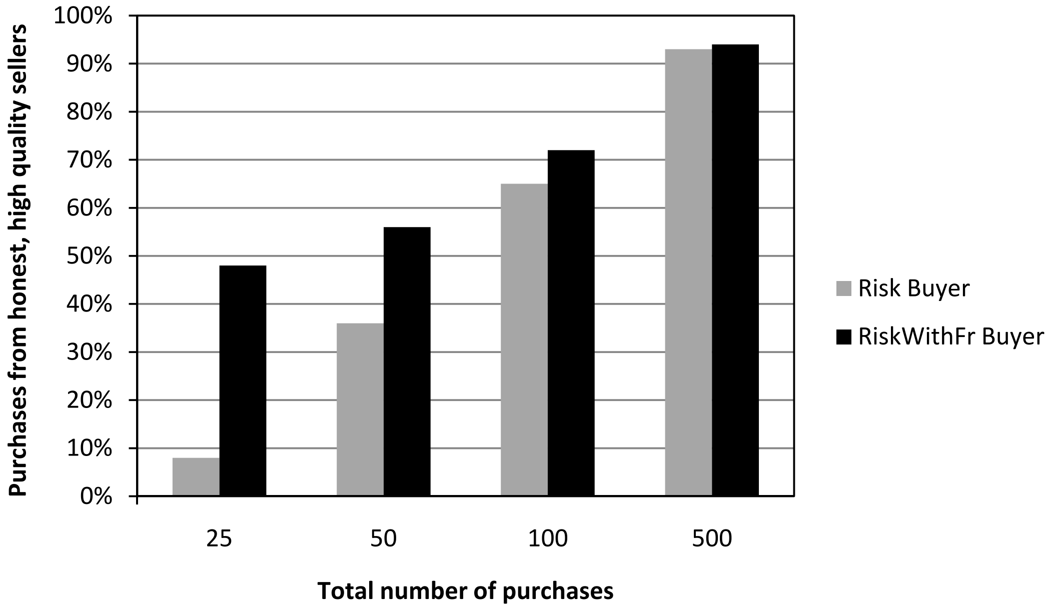

Figure 4.7: Percentage of purchases from honest, high quality sellers by Risk Buyer and RiskWithFr Buyer versus total number of purchases with 32 sellers in the market.

It can be seen in figure 4.7 that at a lower number of purchases, RiskWithFr Buyer's performance is much higher than that of Risk Buyer. Over 100 trials, in 25 purchases, Risk Buyer made on average only 8% of total purchases from honest, high quality sellers as compared to RiskWithFr Buyer who made 48% of total purchases from honest, high quality sellers which is an improvement of 500%. In 50 purchases, the average percentage of purchases from honest, high quality sellers by Risk Buyer was 36%, and by RiskWithFr Buyer was 56%, an improvement of 55%. In 100 transactions, the average percentage of purchases from honest, high quality sellers by Risk Buyer was 65% and RiskWithFr Buyer was 72%, a percentage increase of 11%. In 500 purchases, Risk Buyer made 93% and RiskWithFr Buyer made 94% of total purchases from honest, high quality sellers. When the buyers are making frequent purchases, they are in the market long enough to figure out who the best sellers are and then tend to purchase frequently from them. Hence the difference in the performances of the buyer acting alone and a buyer with friends reduces as the purchase frequency increases.

From the data collected in experiment 2, we compared the average percentage of purchases made from honest, high quality sellers in 50 purchases, by Risk Buyer and RiskWithFr Buyer, as the number of sellers was varied. Figure 4.8 shows the average percentage of purchases made from honest, high quality sellers in 50 purchases by Risk Buyer and RiskWithFr Buyer as the number of sellers was varied.

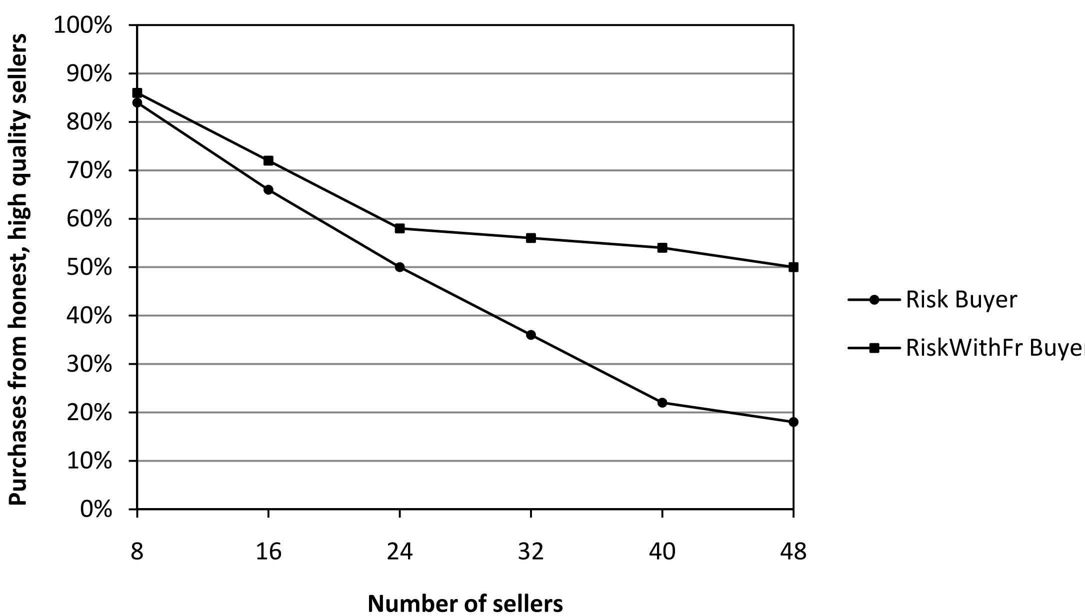

Figure 4.8: Percentage of purchases from honest, high quality sellers by Risk Buyer and RiskWithFr Buyer in 50 transactions as the number of sellers is varied.

From Figure 4.8, we see that when there were 8 sellers, the difference between purchases made from honest, high quality sellers by Risk Buyer (84%) and RiskWithFr (86%) Buyer was marginal. As the number of sellers was increased in steps of 8, RiskWithFr Buyer made a significantly higher number of purchases from honest, high quality sellers than Risk Buyer. When the market was populated by 24 or more sellers, RiskWithFr Buyer made 50 to 58% of total purchases, and Risk Buyer made 18 to 50% of total purchases from honest, high quality sellers, confirming that as the seller count increases, performance of RiskWithFr Buyer is higher.

Our next experiment was to determine to what degree a buyer's opinions should match with a friend's opinions, so that utilizing that friend's opinions helps in improving the buyer's performance. In experiment 3 we kept the seller configuration the same as in experiment 1. We considered six friends who were honest and with a seller rating scale of [-1, 1]. They were different in the percentage of time their seller rating matched that of RiskWithFr Buyer, as described next. We looked at the performances of Risk Buyer and RiskWithFr Buyer and the latter had one the following six friends:

1. *100Acc Friend:* This friend's opinions matched RiskWithFr Buyer's opinions 100 % of the time.

2. *90Acc Friend:* This friend's opinions matched RiskWithFr Buyer's opinions 90% of the time.

3. *80Acc Friend:* This friend's opinions matched RiskWithFr Buyer's opinions 80% of the time.

4. *70Acc Friend:* This friend's opinions matched RiskWithFr Buyer's opinions 70% of the time.

5. *60Acc Friend:* This friend's opinions matched the RiskWithFr Buyer's opinions 60% of the time.

6. *50Acc Friend:* This friend's opinions matched RiskWithFr Buyer's opinions 50% of the time.

We conducted 100 trials for each type of friend, and in each trial the buyers made 100 transactions. We had the friend make its 100 purchases first so that it would provide useful information to RiskWithFr Buyer. We set α=0.5, β=-0.7, as we wanted the friend to be identified early. We looked at how many purchases were made from honest, high quality sellers, and we looked at the percentage of total transactions when a friend's opinions were used when making purchases and this information is shown in Table 4.7.

Table 4.7: Risk Buyer's and RiskWithFr Buyer's performances with friends of different degrees of similarity in opinions

Buyer	Friend	Percentage of purchases from honest high quality sellers in X transactions			Average percentage of transactions a friend's opinion is used in X Transactions		
		X=25	X=50	X=100	X=25	X=50	X=100
Risk Buyer	None	8%	34%	65%	-	-	-
RiskWithFr Buyer	100Acc Friend	68%	80%	89%	92%	94%	97%
RiskWithFr Buyer	90Acc Friend	40%	36%	65%	88%	70%	80%
RiskWithFr Buyer	80Acc Friend	16%	36%	65%	68%	60%	70%
RiskWithFr Buyer	70Acc Friend	12%	36%	66%	60%	46%	71% (In 50% of trials) 15% (In 50% of trials)
RiskWithFr Buyer	60Acc Friend	12%	36%	66%	52%	28%	13% (In 90% of trials) 68% (In10% of trials)
RiskWithFr Buyer	50Acc Friend	12%	36%	65%	32%	18%	9%

From Table 4.7 it can be seen that RiskWithFr Buyer makes the maximum number of purchases from honest, high quality sellers when it has a friend whose opinions match 100% of the time to its own. At lower number of purchases, the performance of RiskWithFr Buyer is substantially higher than Risk Buyer when utilizing the opinions of friends whose opinions match up to 80% of the time, whereas at higher frequency purchases the performance is substantially higher only when the opinions match 100% of the time. As the similarities of a friend's opinions reduce, the percentage time its opinions are utilized is also lower and the performance of RiskWithFr Buyer degrades down to Risk Buyer's performance. In the case of 100 purchases, the percentage time a friend's opinion is utilized decreases significantly for friends whose opinions match less than 70% of the time, as compared to lower frequency purchases. At higher purchase frequency, RiskWithFr Buyer has sufficient interactions with the sellers to form its own impressions to realize that utilizing the opinion of such friends is not beneficial to it.

Next we conducted experiment 4 to determine if there is a range of optimal values for α and β, the incremental and decreasing factors for adjusting a friend's reputation. The seller configuration was the same as in experiments 1 and 3. RiskWithFr Buyer had seven friends out of which six were the same as used in experiment 3, and the seventh friend was the Erratic Friend used in experiments 1 and 2. We set different values of α and β and conducted 100 trials for each combination of α and β. We compared the performance of RiskWithFr Buyer for different values of α and β on the following parameters: percentage of purchases from honest, high quality sellers; number of times each type of friend was identified as trustworthy in 100 trials expressed as a percentage; and average percentage a friend's opinion was utilized when identified as trustworthy in each trial. This data is shown in Tables 4.8 and 4.9.

Table 4. 8 : Comparison of the performance of RiskWithFr Buyer for different values of α, and β

No	α, β	Percentage of purchases from honest, high quality sellers in 100 Purchases	First column under each friend is the average percentage of trials a friend is identified as trustworthy in 100 trials, and the second column shows the average number of times a friend's opinion is used in 100 transactions in each trial					
			100Acc Friend		90Acc Friend		80Acc Friend	
1	.5,-.9	81%	100%	96	100%	86	100%	77
2	.5,-.7	82%	100%	97	100%	89	100%	82
3	.5,.-5	81%	100%	97	100%	92	100%	86
4	.4 -.9	81%	100%	95	100%	82	100%	70
5	.4, -.6	83%	100%	96	100%	87	100%	78
6	.4, -.4	82%	100%	97	100%	90	100%	84
7	.3, -.9	81%	100%	93	100%	72	100%	51
8	.3, -.7	81%	100%	94	100%	81	100%	63
9	.3,-.5	82%	100%	95	100%	83	100%	72
10	.3,-.3	80%	100%	95	100%	88	100%	81
11	.2, -.8	80%	100%	90	100%	60	85%	31
12	.2, -.6	80%	100%	92	100%	68	100%	45
13	.2, -.4	80%	100%	93	100%	76	100%	61
14	.2, -.3	82%	100%	93	100%	81	100	70
15	.2 -.2	80%	100%	94	100%	85	100%	77
16	.1,-.9	76%	100%	70	65%	20	0%	0
17	.1, -0.2	75%	100%	87	100%	69	100%	51

Table 4. 9 :Comparison of the performance of RiskWithFr Buyer for different values of α, and β

No	α, β	First column under each friend is the average percentage of trials a friend is identified as trustworthy in100 trials, and the second column shows the average number of times a friend's opinion is used in 100 transactions in each trial							
		70Acc Friend		Acc60 Friend		Acc50 Friend		Erratic Friend	
1	.5,-.9	100%	61	100%	39	80%	30	35%	6
2	.5,-.7	100%	73	100%	55	93%	47	38%	10
3	.5,.-5	100%	81	100%	73	100%	58	48%	21
4	.4 -.9	100%	46	100%	19	100%	0	25%	1
5	.4, -.6	100%	68	100%	52	68%	38	43%	7
6	.4, -.4	100%	78	100%	70	94%	57	61%	19
7	.3, -.9	100%	29	100%	1	0%	0	14%	5
8	.3, -.7	100%	49	100%	10	0%	0	8%	8
9	.3,-.5	100%	60	100%	35	0%	0	30%	4
10	.3,-.3	100%	75	100%	66	85%	50	59%	18
11	.2, -.8	0%	0	0%	0	0%	0	2%	7
12	.2, -.6	60%	24	0%	0	0%	0	3%	6
13	.2, -.4	99%	44	1%	9	0%	0	10%	10
14	.2, -.3	100	59	92%	35	0%	0	8%	10
15	.2 -.2	100%	70	100%	58	85%	34	47%	16
16	.1,-.9	0%	0	0%	0	0%	0	2%	5
17	.1, -0.2	90%	26	1%	36	1%	30	1%	32

The desirable values for α, β would be when the following criteria are satisfied.

1. Percentage of purchases from honest high quality sellers should be high.

2. Friends whose opinions match up to 80% of the time should be identified as trustworthy 100% of the time, and the number of times their opinion is utilized should be high, as they can be helpful to the buyer when making infrequent purchases as shown in Table 4.7.

3. Friends whose opinions do no match greater than 60% of the time should not be identified as trustworthy or, if identified as such, the number of times their opinions is used should be very low (as we saw in Table 4.7, when the buyer has sufficient time to make judgments about sellers on its own, the number of times the opinions of untrustworthy friends is utilized is very low).

4. Erratic friends should not be identified as trustworthy or, if identified as such, the number of times their opinion is utilized should be very low.

Applying the above criteria to the data collected in experiment 4, we narrowed down the preferred values for α, β and these values and the performance of RiskWithFr Buyer at these values are highlighted in tables 4.8 and 4.9. Our results indicate that $0.2 \leq \alpha \leq 0.4$ and $2\alpha \leq \beta \leq 3\alpha$. For α less than 0.2, the average percentage of purchases from honest, high quality sellers is lower than when α is greater than or equal to 0.2 (lines 16 and 17 in table 4.8). When $\alpha \geq 0.4$ or $\beta=\alpha$, friends whose opinions match less than 70% of the time are identified frequently and number of times their opinions is utilized is high (lines 1, 2, 3, 10 and 15 in Table 4.9). When β is slightly greater than α, then friends whose opinions match 60% of the time are identified frequently and their opinions utilization is high (lines 9 and 14 in Table 4.9). When the value of $\beta \geq 3\alpha$ friends whose opinions match 80% of time or greater are not identified as trustworthy all the time and their opinions utilization is low (line 11 in Table 4.8). Thus α should be between 0.2 and 0.4, and for a given value of α, β should be 2 to 3 times the value of α.

4.4 Phase III

In this phase, we compared the two different methods of identifying trustworthy friends and utilizing the information provided by them. Data was collected from experiments 1 and 2 as described in section 4.2. In experiment 1, we populated the market with 32 sellers, with four sellers from each category in table 4.3. We also populated the market with 14 buyers comprising of 6 buyers, one from each category in table 4.4 and one of each type of 8 friends summarized in table 4.5. We conducted 100 trials. In each trial buyers made 500 purchases or transactions with sellers. In experiment 2, the buyer configuration was the same as in experiment 1. We varied the number of sellers in the market and conducted 100 trials for each seller count. The buyers made 500 transactions in each trial. In this phase, we looked at the performances of buyers when making 50 purchases only, as the goal was to study how different buyers performed when making fewer purchases with different sellers in the market. We analyzed the buyers' performance with different seller counts at different purchase levels in phase V.

From the data collected in experiments 1 and 2, performances of RiskWithFr Buyer and OOnlyRiskWithFr Buyer were compared. RiskWithFr Buyer uses the first method (section 3.2.1) of identifying trustworthy friends. In this method, 1) friends' opinions and numerical ratings of sellers are used to identify trustworthy friends, and, 2) trustworthy friends' opinions and ratings of sellers along with the RiskWithFr Buyer's own information are used to choose a seller to purchase from.

OOnlyRiskWithFr Buyer uses the second method (section 3.2.2) of identifying trustworthy friends. In this method, 1) only friends' opinions of sellers are used to identify trustworthy friends, and, 2) trustworthy friends' opinions of sellers along with OOnlyRiskWithFr Buyer's own information are used to choose a seller to purchase from. The two buyers, RiskWithFr Buyer and OOnlyRiskWithFr Buyer, were compared in their abilities: to identify trustworthy friends, time taken to identify trustworthy friends, and the percentage of purchases from honest, high quality sellers at different levels of purchase, and as the number of sellers in the market was varied. We wanted to determine which method (first or second) of identifying trusted friends and using information from them is superior.

Table 4.9 shows the performance of the two buyers against different friend types.

Table 4.10: Comparison of RiskWithFr (RFr) Buyer and OOnlyRiskWithFr (OORFr) Buyer against different friends

Friend	Average transaction when a friend is first identified as a trustworthy friend		Average percentage of 500 transactions a friend's opinion is used in each trial.		Percentage of 100 trials identified as a trustworthy friend	
	RFr Buyer	OORFr Buyer	RFr Buyer	OORFr Buyer	RFr Buyer	OORFr Buyer
Friends with similar opinions						
Risk Friend	13	6	96%	99%	100%	100%
CD-0.5 Friend	8	6	96%	99%	100%	100%
CM-10 Friend	39	6	76%	99%	100%	100%
N-01 Friend	11	6	94%	99%	100%	100%
Friends with slightly different opinions						
SDO Friend	28	7	84%	99%	100%	100%
Tran Friend	119	47	73%	91%	100%	100%
Friends with very different opinions						
VDO Friend	38	6	3%	0.2%	14%	11%
Lying Friend						
Erratic Friend	20	8	0.4%	4%	3 %	7%

The desired behavior from the two buyers is that honest friends with similar opinions (Risk Friend, CD-0.5 Friend, CM-10 Friend, and N-01 Friend) are identified as trustworthy at the earliest and their opinions are utilized all the time once they are identified. Friends with very different opinions (VDO Friend) or lying friends (Erratic Friend) should not be identified as a trustworthy friend at all, or if erroneously identified, buyers should be able to quickly realize their mistake and not utilize the opinions of such friends. As it would take longer to identify friends with slightly different opinions (SDO Friend and Tran Friend) as trustworthy, consequently, the opinions of such friends should be utilized a lesser number of times as compared to honest friends with similar opinions.

From table 4.9 it can be seen that, with regard to honest friends with similar opinions, OOnlyRiskWithFr Buyer identifies them earlier than RiskWithFr Buyer and also utilizes their opinions a larger number of times. For friends whose opinions differed slightly (SDO Friend and Tran Friend), OOnlyRiskWithFr Buyer identified them as trustworthy a lot earlier than RiskWithFr Buyer and also utilized their opinions to a higher extent than RiskWithFr Buyer. VDO Friend which had very different opinions was erroneously identified as trustworthy in 14% of the 100 trials by RiskWithFr Buyer and in 11% of the trials by OOnlyRiskWithFr Buyer. Both buyers are able to recognize their mistake, which is indicated by the fact that VDO Friend's opinions were utilized on an average of 3% by RiskWithFr Buyer and 0.2% by OOnlyRiskWithFr Buyer. Erratic Friend was incorrectly identified as a trustworthy friend in 3% of the 100 trials by RiskWithFr Buyer and its opinions were utilized less than half percent of all transactions. Erratic friend was incorrectly identified as a trustworthy friend in 8% of the 100 trials by OOnlyRiskWithFr Buyer and its opinions were utilized on an average of 4% of all transactions.

Our results indicate that compared to RiskWithFr Buyer, OOnlyRiskWithFr Buyer does well in good to optimal conditions. However, it does a lot worse in adversarial environment. When the friends were honest in their replies, irrespective of their similarities in opinions, performance of OOnlyRiskWithFr Buyer's performance was superior to RiskWithFr Buyer's performance. However, in the case of lying friend, RiskWithFr Buyer's performance was better. It erroneously identified lying friend as trustworthy in fewer trials, and when the lying friend was incorrectly identified as trustworthy, RiskWithFr Buyer was able to correct its mistake a lot sooner than OOnlyRiskWithFr.

Next we compared the performance of the two buyers on the number of purchases made from honest, high quality sellers. Figure 4.9 shows the percentage of purchases made from honest, high quality sellers by RiskWithFr Buyer and OOnlyRiskWithFr Buyer at different levels of purchases.

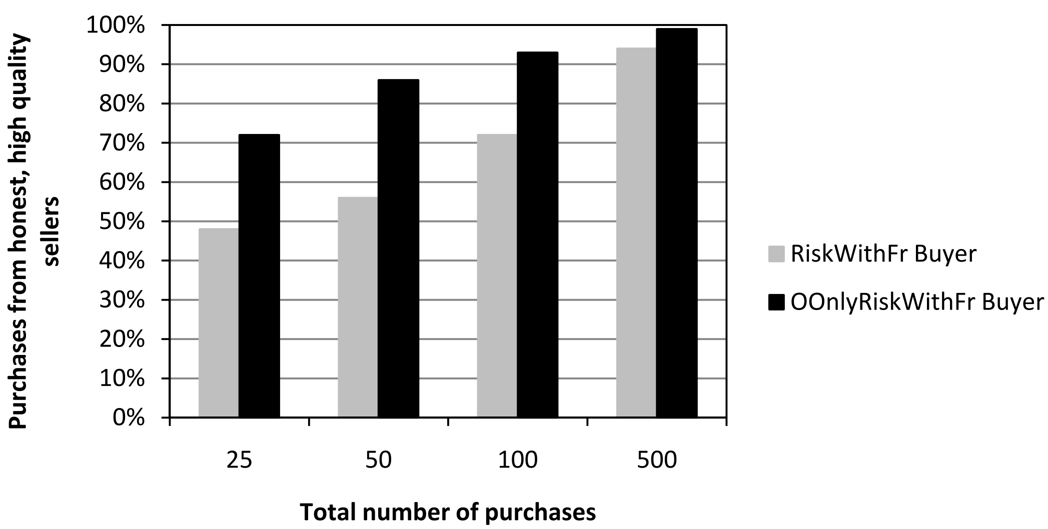

Figure 4.9: Percentage of purchases from honest, high quality sellers by RiskWithFr Buyer and OOnlyRiskWithFr Buyer versus total number of purchases with 32 sellers in the market.

It can be seen that, OOnlyRiskWithFr Buyer outperforms RiskWithFr Buyer at all levels of purchases, and the percentage improvement is higher at a lower number of purchases. Over 100 trials, in 25 purchases, RiskWithFr Buyer made on average 48% of total purchases from honest, high quality sellers as compared to OOnlyRiskWithFr Buyer who made 72% of total purchases from honest, high quality sellers, which is an improvement of 50%. In 50 purchases, the average percentage of purchases from honest, high quality sellers by RiskWithFR Buyer was 56%, and the OOnlyRiskWithFr Buyer made 86% of total purchases, an improvement of 53%. In 100 transactions, the average percentage of purchases from honest, high quality sellers by RiskWithFr Buyer was 72% and by OOnlyRiskWithFr Buyer was 93%, a percentage increase of 29%. In 500 purchases, the average percentage of purchases from honest, high quality sellers by RiskWithFr Buyer was 94% and OOnlyRiskWithFr Buyer made 99% of total purchases, an improvement of 5%. When buyers are making frequent purchases, they are in the market long enough to figure out who the best sellers are and then tend to purchase frequently from them. Hence the difference in the performances of RiskBuyerWithFr Buyer and OOnlyRiskBuyerWithFr Buyer reduces as the purchase frequency increases.

From the data collected in experiment 2, we compared the average percentage of purchases made from honest, high quality sellers in 50 purchases, by RiskWithFr Buyer and OOnlyRiskWithFr Buyer, as the number of sellers was varied. Figure 4.10 shows the average percentage of purchases made from honest, high quality sellers in

50 purchases by RiskWithFr Buyer and OOnlyRiskWithFr Buyer as the number of sellers was varied.

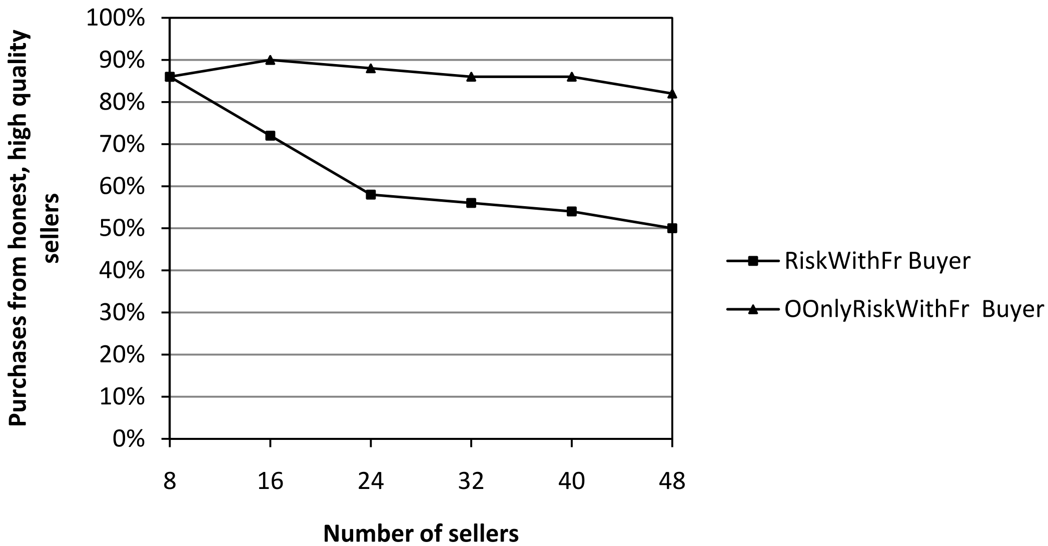

Figure 4.10: Percentage of purchases from honest, high quality sellers by RiskWithFr Buyer and OOnlyRiskWithFr Buyer in 50 purchases as the number of sellers is varied.

It can be seen from figure 4.10, that as the number of sellers is increased, there is no significant difference in the number of purchases made from honest, high quality sellers by OOnlyRiskWithFr Buyer, whereas the number of purchases made from honest, high quality sellers by RiskWithFr Buyer reduces. From Figure 4.10, we see that when there were 8 sellers, there was no difference between the purchases made from honest, high quality sellers by RiskWithFr Buyer (86%) and OOnlyRiskWithFr Buyer (86%). As the number of sellers was increased in steps of 8, the percentage of purchases from honest, high quality sellers by RiskWithFr Buyer comes down from 86% to 50% whereas the percentage of purchases from honest, high quality sellers by

OOnlyRiskWithFr Buyer remains steady over 80%, confirming that at higher seller counts the performance of OOnlyRskWithFr Buyer is superior.

Overall, OOnlyRiskWithFr Buyer has a superior performance as compared to RiskWithFr Buyer in environments where majority of friends are honest. OOnlyRiskWithFr Buyer is able to utilize the opinions of trustworthy friends more often than RiskWithFr Buyer. Like RiskWithFr Buyer, it is successful in avoiding friends with very different opinions by either identifying them correctly as having different opinions, or if a mistake has been made, it is able to correct it quickly. It is able to make higher number of purchases from honest, high quality sellers than RiskWithFr Buyer at lower frequencies of purchase, and its performance remains fairly constant as the number of sellers is varied, whereas the performance of RiskWithFr Buyer decreases as the number of sellers is increased.

We think OOnlyRiskWithFr Buyer's performance is superior because it is using only friends' opinions in identifying trustworthy friends. RiskWithFr Buyer on the other hand uses friends' opinions and ratings of sellers. RiskWithFr Buyer considers a friend trustworthy when the friend's reputation exceeds a threshold value (which happens because its opinions regarding sellers have been similar to RiskWithFr Buyer), and the standard deviation of the differences in RiskWithFr Buyer's ratings and the friend's ratings of sellers is below a certain threshold value. So if the opinions continue to match, but the standard deviation of the differences between RiskWithFr Buyer's ratings and the friend's ratings of sellers falls below the threshold value, then that friend will not be considered trustworthy and that friend's information will not be incorporated into the decision making. There will be a consistent difference between the buyer's rating and a friend's rating only if the rating mechanisms of the buyer and the friend follow the same curve. Friends may be honest and their recommendation for a seller may be similar to the buyer's, however if the standard deviation of the differences in ratings is not small, then their opinions will not be considered in the first approach. On the other hand, OOnlyRiskWithFr Buyer considers a friend as trustworthy if the friend's reputation simply exceeds a threshold value. As long as the friend's opinions and OOnlyRiskWithFr Buyer's opinions match, then that friend is included in the list of trustworthy friends whose information should be used in decision making.

Our results also show that OOnlyRiskWithFr Buyer's performance is inferior to RiskWithFr Buyer in detecting the nature of lying friends, and leads us to conclude that its performance will be worse than RiskWithFr Buyer in adversarial conditions where majority of friends are either erratic or deliberately lie. While utilizing the opinions and rating of sellers by friends to detect trustworthy friends is a slower cautious approach, it also enables the buyer to correctly detect untrustworthy friends. Utilizing only opinions of sellers to identify trustworthy friends is faster, but it also makes the buyer more vulnerable in adverse conditions.

4.5 Phase IV

In this phase we compared conservative and risk taking attitudes of buyers. We wanted to see how attitude affects the performance of the buyers when they are acting alone and when they have friends.

From the data collected in experiment 1 and 2 (section 4.2), first, we compared the performances of Risk Buyer and Cns Buyer. Risk Buyer and Cns Buyer make decisions regarding purchases based on their own experiences with sellers and use our direct trust rating model as described in section 3.1 in chapter 3 for evaluating sellers. A Risk Buyer's attitude towards new buyers is "Risk Taking," which means that new or unexplored buyers are always included in the list of potential sellers. Cns Buyer purchases from new sellers during its exploration of the market, where the exploration rate is proportional to the ratio of new sellers to all the sellers in the market. The detailed policy of a buyer with conservative attitude regarding new sellers is described in section 3.5 of chapter 3. Figure 4.11 shows the percentage of purchases made from honest, high quality sellers by Risk Buyer and Cns Buyer at different levels of purchases when the market was populated with 32 sellers.

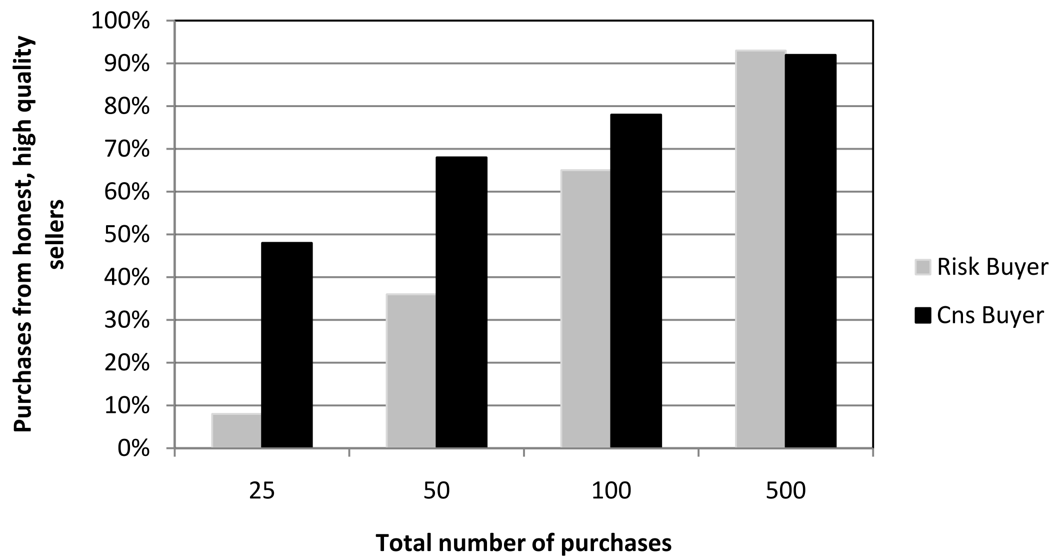

Figure 4.11: Percentage of purchases from honest, high quality sellers by Risk Buyer and Cns Buyer versus total number of purchases with 32 sellers in the market.

It can be seen that Cns Buyer has a superior performance compared to Risk Buyer at a lower number of purchases. Over 100 trials, in 25 purchases, Risk Buyer made on average 8% of its total purchases from honest, high quality sellers as compared to Cns Buyer who made 48% of total purchases from honest, high quality sellers, which is an improvement of 500%. In 50 purchases, the average percentage of purchases from honest, high quality sellers by Risk Buyer was 36%, and Cns Buyer made 68% of total purchases, an improvement of 89%. In 100 purchases, the average percentage of purchases from honest, high quality sellers by Risk Buyer was 65% and by Cns Buyer was 78%, a percentage increase of 20%. In 500 purchases, the average percentage of purchases from honest, high quality seller by Risk Buyer was 93% and Cns Buyer made 92% of total purchases, a decrease of 1.1%. When buyers are making frequent purchases, they are in the market long enough to figure out who the best sellers are and then tend to purchase frequently from them. Hence the difference in the performances of Risk Buyer and Cns Buyer reduces at high levels of purchase.

From data collected in experiment 2, we compared the average percentage of purchases made from honest, high quality sellers in 50 purchases, by Risk Buyer and Cns Buyer, as the number of sellers was varied. This is shown in figure 4.12.

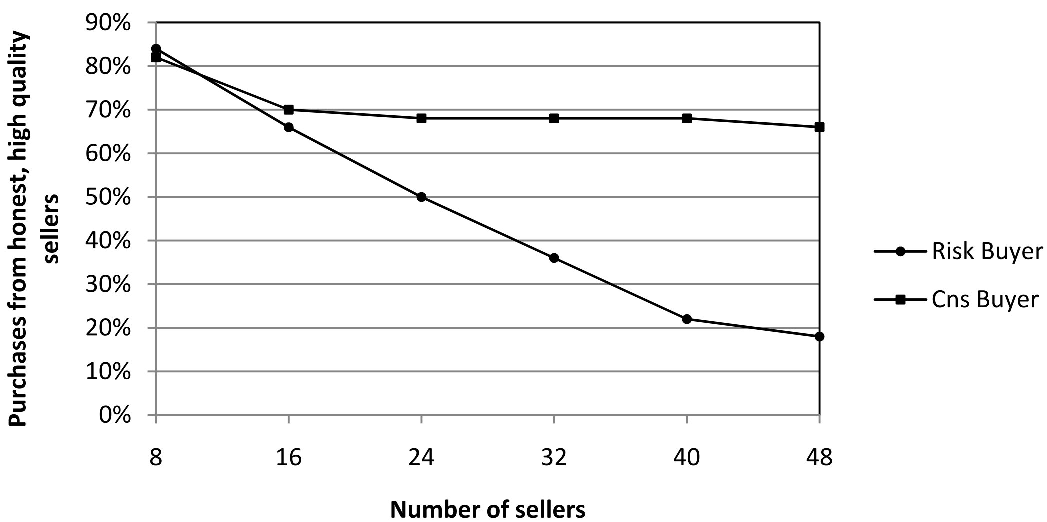

Figure 4.12: Percentage of purchases made by Risk Buyer and Cns Buyer from honest, high quality sellers in 50 purchases as the number of sellers is varied.

From Figure 4.12, we see that when there were up to 16 sellers in the market, the differences between Risk Buyer (84% for 8 sellers and 66% for 16 sellers) and Cns Buyer (82% for 8 sellers and 70% for 16 sellers) were marginal. As the number of sellers was increased in steps of 8, the percentage of purchases from honest, high quality sellers by Risk Buyer came down from 66% to 18%, whereas the percentage of purchases from honest, high quality sellers by Cns Buyer remained steadily over 65%.

When the buyers are acting alone, a buyer with conservative attitude has a superior performance compared to a buyer with risk taking attitude, at lower number of purchases, and at higher numbers of sellers. We think this is because a buyer with risk taking attitude always purchases from new sellers if they are offering a lower price than the best seller that the buyer has purchased from in the past, whereas a conservative buyer is cautious in its approach.

A buyer with conservative attitude had better performance than a buyer with risk taking attitude when acting alone. We wanted to see if this would remain the same or change in the case of buyers with friends. Figure 4.13 shows the average percentage of purchases from honest, high quality sellers by RiskWithFr Buyer and CnsWithFr Buyer at different levels of purchases. Both these buyers use the first method of identifying trustworthy friends (section 3.2.1).

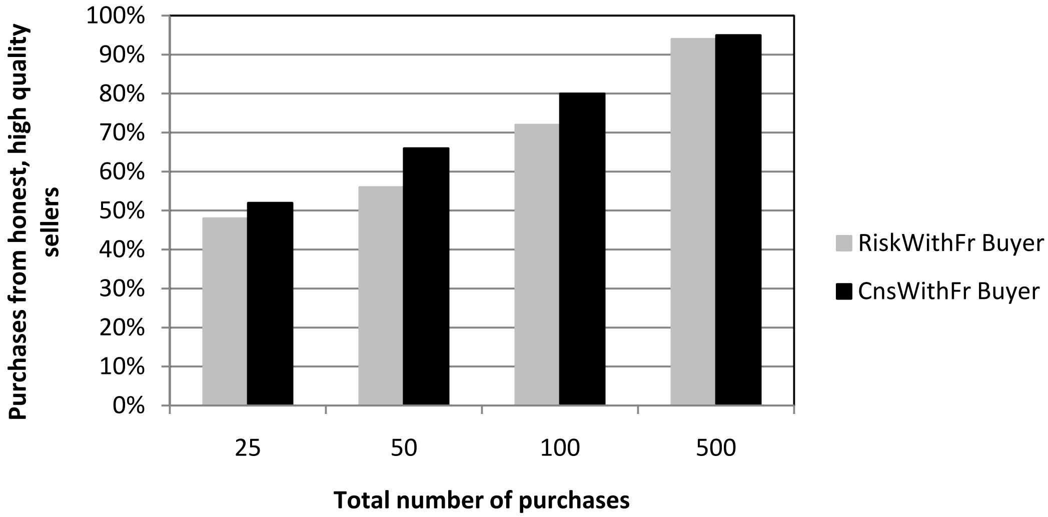

Figure 4.13: Percentage of purchases from honest, high quality sellers by RiskWithFr Buyer and CnsWithFr Buyer versus total number of purchases, when the market is populated with 32 sellers.

It can be seen that, CnsWithFr Buyer's performance is marginally higher than RiskWithFr Buyer's at lower number of purchases. Over 100 trials, in 25 purchases, RiskWithFr Buyer made on average 48% of total purchases from honest, high quality sellers as compared to CnsWithFr Buyer who made 52% of total purchases from honest, high quality sellers, which is an improvement of 8%. In 50 purchases, the average percentage of purchases from honest, high quality sellers by RiskWithFr Buyer was 56%, and CnsWithFr Buyer made 66% of total purchases, an improvement of 18%. In 100 transactions, the average percentage of purchases from honest, high quality sellers by Risk WithFr Buyer was 72% and by CnsWithFr Buyer was 80%, a percentage increase of 11%. In 500 purchases, the average percentage of purchases from honest, high quality seller by RiskWithFr Buyer was 94% and CnsWithFr Buyer made 95% of total purchases, an increase of 1%.

From data collected in experiment 2, we compared the average percentage of purchases made from honest, high quality sellers in 50 purchases, by RiskWithFr Buyer and CnsWithFr Buyer, as the number of sellers was varied. This is shown in figure 4.14.

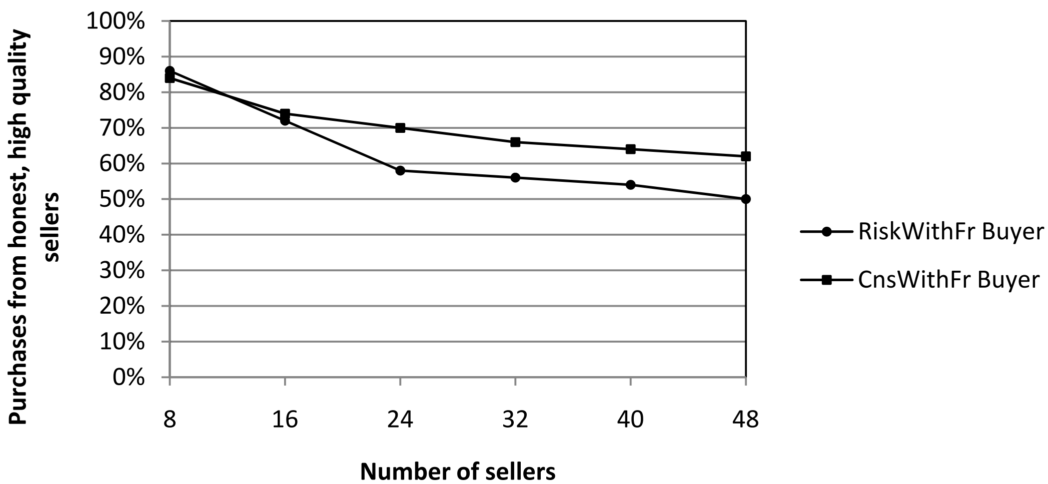

Figure 4.14: Percentage of purchases made from honest, high quality sellers by RiskWithFr Buyer and CnsWithFr Buyers in 50 purchases as the number of sellers is varied.

From Figure 4.14, we see that as the number of sellers was increased, the average percentage of purchases from honest, high quality sellers came down for both sellers. As the number of sellers was increased in steps of 8, the percentage of purchases from honest, high quality sellers by RiskWithFr Buyer came down from 86% to 50%, whereas the percentage of purchases from honest, high quality sellers by CnsWithFr Buyer came down from 84% to 62%. When the market was populated with more than 16 sellers, CnsWithFr Buyer's purchases from honest, high quality sellers were 10% higher than RskWithFr Buyer's purchases from the same category of sellers.

When the market had less than 16 sellers, the differences in purchases from honest, high quality sellers by both buyers were marginal.

When buyers used the first method to identify trusted friends, a buyer with conservative attitude, CnsBuyerWithFr Buyer, performed slightly better (by about 8-10%) than a buyer with risk taking attitude, RskWithFr Buyer, at lower levels of purchases and at higher number of sellers.

Figure 4.15 shows the average percentage of purchases from honest, high quality sellers by OOnlyRiskWithFr Buyer and OOnlyCnsWithFr Buyer at different levels of purchases when the market was populated with 32 sellers. Both these buyers use the second method of identifying trustworthy friends (section 3.2.2).

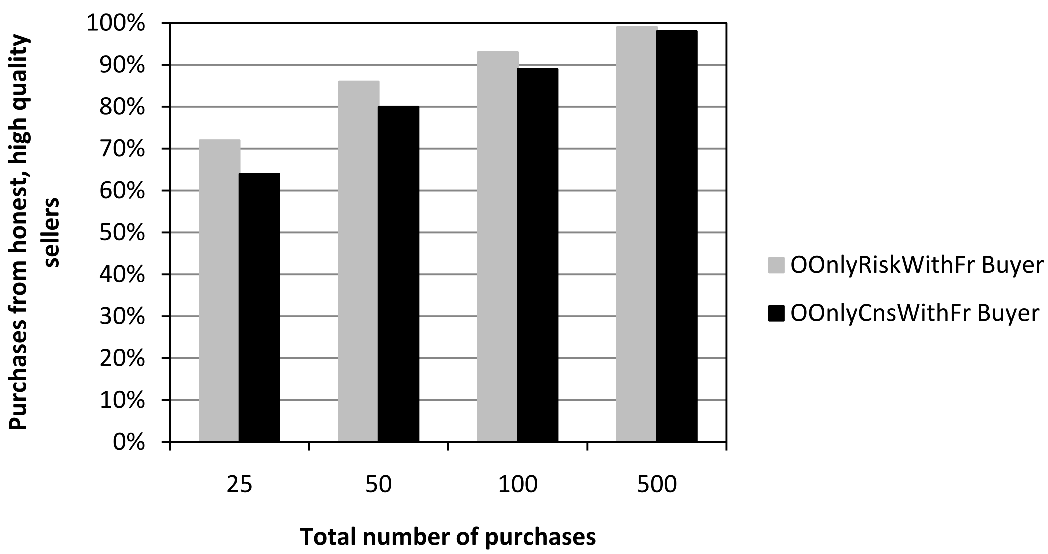

Figure 4.15: Percentage of purchases from honest, high quality sellers by OOnlyRiskWithFr Buyer and OOnlyCnsWithFr Buyer versus total number of purchases when the market is populated with 32 sellers.

It can be seen from figure 4.15 that, OOnlyRskWithFr Buyer's average percentage of purchases from honest, high quality sellers (72% in 25, 86% in 50, 93% in 100, and 99% in 500) is almost the same or marginally higher than OOnlyCnsWithFr Buyer's average percentage of purchases from honest, high quality sellers (64% in 25, 80% in 50, 89% in 100, and 98% in 500).

Figure 4.16 shows the average percentage of purchases from honest, high quality sellers by OOnlyRiskWithFr Buyer and OOnlyCnsWithFr Buyer in 50 purchases as the number of sellers was varied.

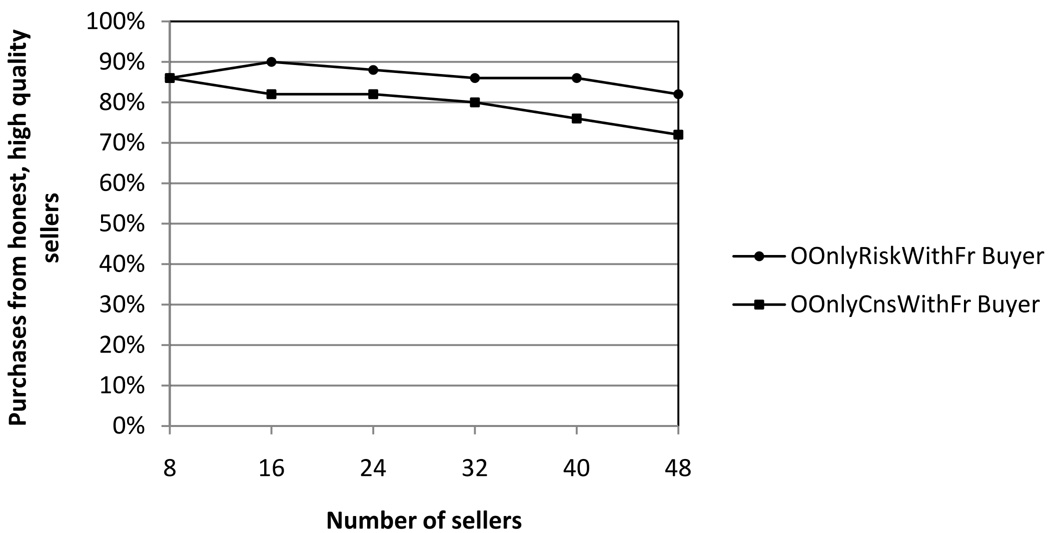

Figure 4.16: Percentage of purchases made from honest, high quality sellers by OOnlyRiskWithFr Buyer and OOnlyCnsWithFr Buyer in 50 purchases as the number of sellers is varied.

From Figure 4.16, it can be seen that as the number of sellers is increased, the average percentage of purchases from honest, high quality sellers remains almost the same for OOnlyRiskWithFr Buyer, and it reduces by about 10% for OOnlyCnsWithFrBuyer.

When buyers used the second method to identify trusted friends, performance of a buyer with risk taking attitude (OOnlyRiskWithFr Buyer) was slightly better than the performance of a buyer with conservative attitude (OOnlyCnsWithFr Buyer) at lower number of purchases and higher number of sellers. On the other hand, when buyers were acting alone or using the first method of identifying trustworthy friends, performance of buyers with conservative attitude was better than buyers with risk taking attitude at lower number of purchases, and at higher number of sellers.

4.6 Phase V

In this phase we compared the performances of all buyers at different levels of purchases, as the number of sellers was varied. We wanted to see if any particular combination of buyer attitude and seller modeling would yield the best performance. From data collected in experiment 2 (section 4.2), we compared the average percentage of purchases made from honest, high quality sellers in 25, 50, 100 and 500 purchases, by Risk Buyer, Cns Buyer, RiskWithFr Buyer, CnsWithFr Buyer, OOnlyRiskWithFr Buyer, and OOnlyCnsWithFr Buyer, as the number of sellers was varied. Figures 4.17 4.18, 4.19, and 4.20 show the percentage of purchases made by various buyers from honest, high quality sellers in 25, 50, 100, and 500 purchases as the number of sellers was varied.

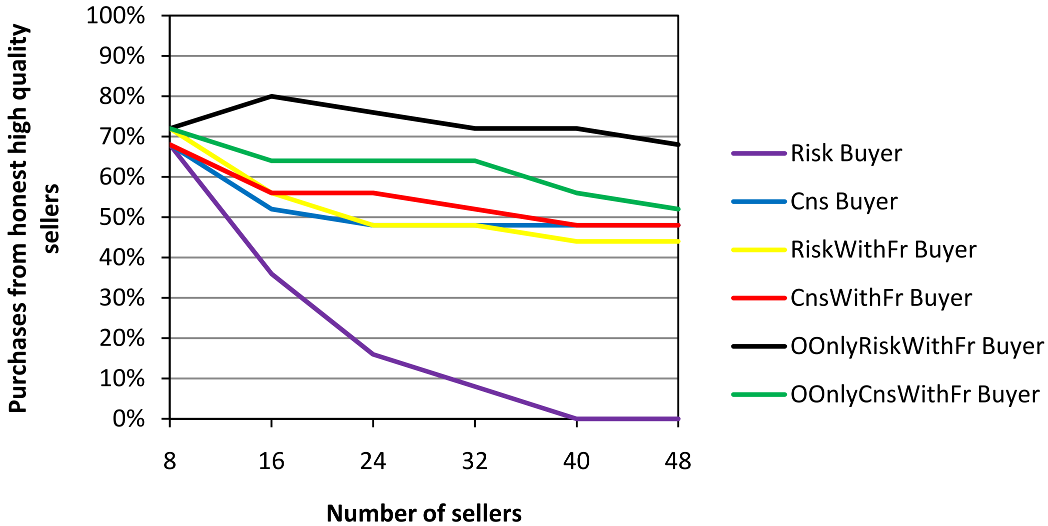

Figure 4.17: Percentage of purchases from honest, high quality sellers by various buyers in 25 purchases as the number of sellers is varied.

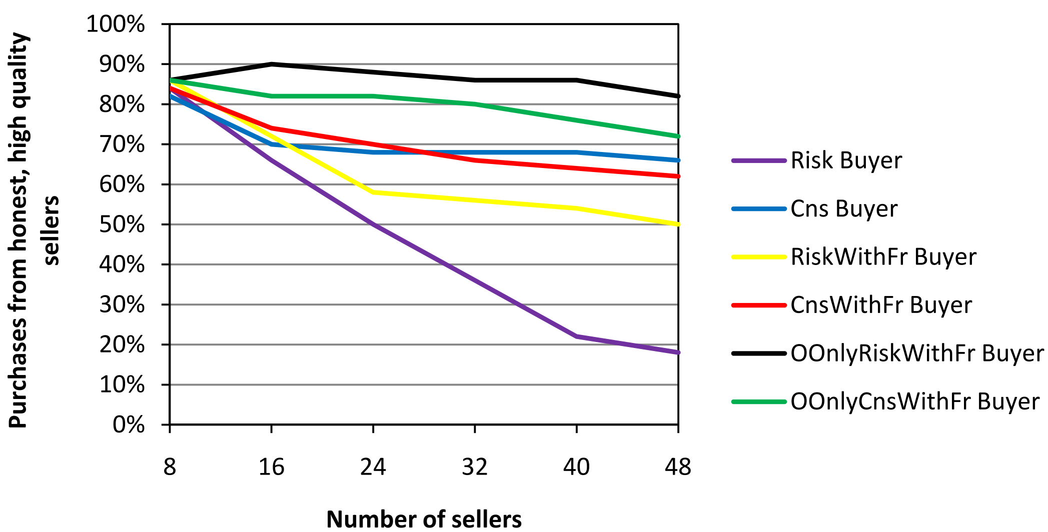

Figure 4.18: Percentage of purchases from honest, high quality sellers by various buyers in 50 purchases as the number of sellers is varied.

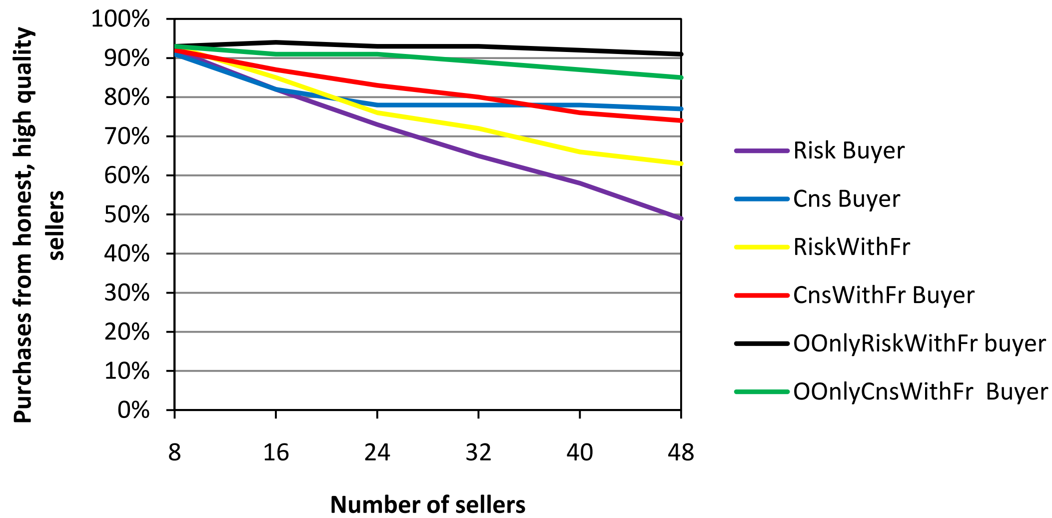

Figure 4.19: Percentage of purchases from honest, high quality sellers by various buyers in 100 purchases as the number of sellers is varied.

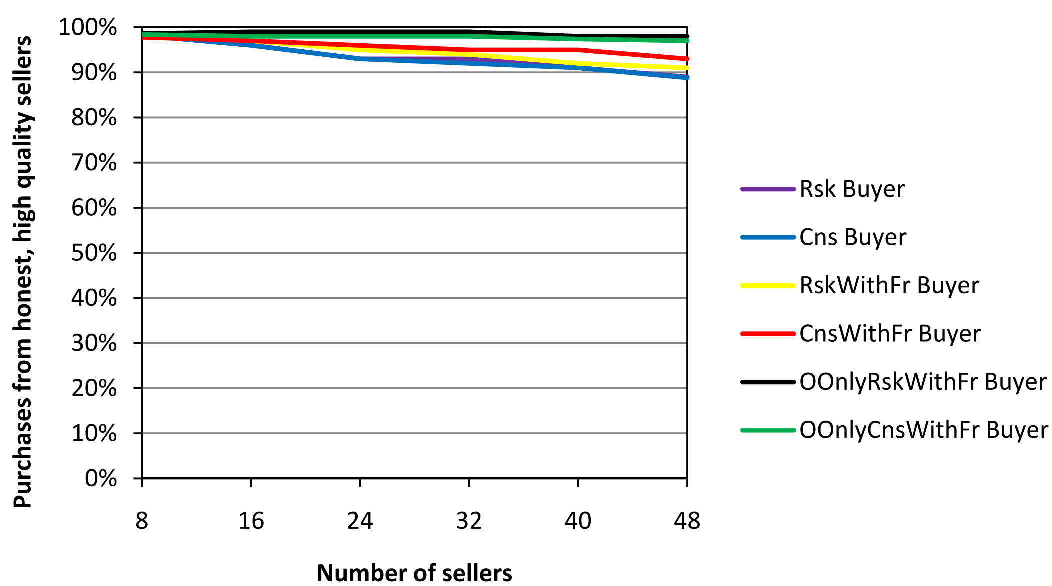

Figure 4.20: Percentage of purchases from honest, high quality sellers by various buyers in 500 purchases as the number of sellers is varied.

Figure 4.20 shows the performance of buyers when making frequent purchases. It can be seen that when there were fewer sellers in the market, the difference in the performance of various buyers was marginal. As the number of sellers was increased, the performances of OOnlyRiskWithFr Buyer and OOnlyCnsWithFr Buyer remained steady; RiskWithFr Buyer and CnsWithFr Buyer experienced a decrease of 5-7% in their performance, and Risk and Cns Buyer experienced a decrease of about 10% in their performance.

Figures 4.17, 4.18 and 4.19 show the performance of buyers when making infrequent purchases. It can be seen that OOnlyRiskWithFr Buyer and OOnlyCnsWithFr Buyer had the best performance. As the number of sellers was increased, OOnlyRiskWithFr Buyer experienced a slight decrease in its performance, whereas, OOnlyCnsWithFr Buyer experienced a higher decrease in its performance. Risk Buyer had the lowest performance among all buyers. Buyers with risk taking attitudes showed significant improvements with both methods of identifying trustworthy friends. RiskWithFr Buyer's performance was higher than Risk Buyer, but lower than OOnlyRiskWithFr Buyer. The performances of RiskWithFr Buyer, CnsWithFr Buyer and Cns Buyer were similar when there were fewer sellers in the market. CnsWithFr Buyer's and Cns Buyer's performances were better than RiskWithFr Buyer's when the number of sellers was high in the market. CnsWithFr Buyer's performance did not show any significant improvement over Cns Buyer performance. Its performance was marginally higher than Cns Buyer's performance at lower number of sellers, and then marginally lower when the number of sellers was high in the market.

OOnlyCnsWithFr Buyer's performance showed a significant improvement over CnsWithFr Buyer and Cns Buyer.

When buyers are making decisions based on their own information, Cns Buyer, a buyer with conservative attitude has the best performance. Intuitively this makes sense. When a buyer is acting alone, the buyer needs to be prudent.

When buyers are utilizing information provided by their trusted friends, OOnlyRiskWithFR Buyer, a buyer with risk taking attitude and using the second method of identifying trusted friends had the best performance. Buyers with conservative and risk taking attitudes utilizing only the trusted friends' opinions (OOnlyRiskWithFr Buyer and OOnlyCnsWithFr Buyer) fared better than buyers using trusted friends' opinions and ratings of sellers (RiskWithFr Buyer and CnsWithFrBuyer).

4.7 Summary

To test our model, we developed a simulation of an electronic market consisting of a Matchmaker [31], buyer agents, and seller agents, which communicate with each other. At any time Matchmaker has a current list of all sellers active in the market. Buyers obtain list of sellers for a product from the Matchmaker. Next, they obtain price quotes from sellers by contacting them individually. Buyers may belong to a community comprising of buyers. Members of the community are referred to as *friends*. While evaluating sellers for purchase decisions a buyer, with friends, requests for seller information from friends. However, there is no guarantee that friends are honest or sharing similar opinions. In our model:

- We presented a method for a buyer to evaluate sellers based on its direct interactions.
- We presented two methods of identifying trustworthy (honest and share similar opinions) friends and utilizing information from them.
- We considered risk taking and conservative attitude of buyers.

We tested various components of our model in five phases. In Phase I, our goal was to test buyer's method of evaluating sellers based on direct interactions. We populated the market with various sellers (varying quality, honest, dishonest, and varying prices) and compared the performance of a buyer using our method of rating sellers based on direct interactions with buyers using models, for rating sellers based on direct interactions, proposed by other researchers (Tran [49], and Vidal and Durfee [53]). The time taken by various buyers to learn to identify high quality, low priced sellers was used as a metric to compare various buyers. In our experiment each buyer made 500 transactions and in each transaction they purchased from a seller using their learning strategy. Our results show that all buyers learn at different rates to identify high quality, low priced sellers. After having learnt, they consistently interact with high quality, low priced sellers. A buyer using our method of rating sellers based on direct interactions was the quickest to learn in 15 transactions. A buyer using Tran's model [49] took about 60 transactions and a buyer using Vidal and Durfee's model [53] took about 250 transactions.

Our experimental setup for phases II to V was the same. We populated the market with various sellers (varying quality, honest, dishonest and varying prices) and with different types of buyer friends (honest, lying, similar opinions, different opinions). We studied the performances of various buyers with conservative attitude or risk taking attitude, acting alone using our method of rating sellers based on direct interactions, or with friends using our method of rating sellers based on direct interactions and one of the two methods of identifying trustworthy (honest and sharing similar opinions) friends that has been presented in this dissertation. All the buyers made 500 transactions and in each transaction they obtained seller list from Matchmaker, obtained price quotes from sellers, requested for seller information from friends if they had friends, waited for a certain period of time for responses and made their decision to purchase from a seller. We had the friends make their purchases first so that they could provide useful information to buyers requesting for it.

In Phase II, our goal was to test the first method of identifying and utilizing information from trustworthy friends as described in section 3.2.1. In this method, the buyer requests for friend's opinions, and ratings of sellers. The buyer uses friend's opinions and standard deviation of the differences between friends' ratings and it's ratings of sellers to identify trustworthy friends. The reputation ratings of sellers provided by trustworthy friends are adjusted to account for the differences in the friends' and the buyer's rating systems and combined with the buyer's rating of sellers to identify potential sellers to purchase from. The metrics used were: ability to identify trustworthy friends, time taken to identify trustworthy friends, and percentage of purchases from honest, high quality sellers at different levels of purchase, and as the number of sellers in the market was varied. Desired behavior is that honest friends who share similar opinions are identified at the earliest and their information about sellers utilized all the time. Friends who are not honest, or who do not share similar opinions should not be identified as trustworthy and their information regarding sellers should not be utilized at all. Buyer with friends should make higher percentage of purchases than buyer acting alone from honest, high quality sellers. Our results show that honest friends with similar opinions were identified as trustworthy in all the 100 trials, their opinions were utilized from 76% to 96% of 500 transactions in each trial. Honest friends with slightly different opinions were also identified as trustworthy in all the 100 trials, and their opinions were utilized from 72% to 84% of 500 transactions. Honest friend with very different opinions was successfully identified as not trustworthy in 86 of the 100 trials. In 14 trials, it was erroneously identified as trustworthy. However the buyer was able to realize its mistake which is indicated by the fact that its opinions were utilized only 3% of 500 transactions in each of the 14 trials. Lying friend was correctly identified as not trustworthy in 97 of 100 trials. It was incorrectly identified as trustworthy in 3 trials. However its opinions were utilized less than 1% of 500 transactions in each of those 3 trials. Our results also show that once trustworthy friends have been identified and their opinions have been utilized in decision making, a buyer using our first model of identifying trustworthy friends has superior performance than a buyer acting alone at lower levels of purchases and for increasing numbers of sellers in the market.

In Phase III, we evaluated the second method of identifying and utilizing information from trustworthy friends as described in section 3.2.2. In this method, the buyer requests for only friend's opinions of sellers. The buyer uses friend's opinions to identify trustworthy friends. Based on a trusted friend's opinion, reputation value is assigned to a seller. Assigned reputation values and the buyer's ratings of sellers are combined to identify potential sellers to purchase from. We compared the performances of buyers with friends using first and second method of identifying trustworthy friends. We wanted to determine if any method (first or second) of identifying trusted friends and using information from them is superior. Metrics used in Phase III were exactly the same as in Phase II.

Results for a buyer using the first method of identifying trustworthy friends were presented in Phase II. Our results show that for a buyer using the second method, honest friends with similar opinions were identified as trustworthy in all the 100 trials, and their opinions were utilized in 99% of 500 transactions in each trial. Honest friends with slightly different opinions were also identified as trustworthy in all the 100 trials, and their opinions were utilized from 91% to 99% of 500 transactions in each trial. Honest friend with very different opinions was successfully identified as not trustworthy in 89 of the 100 trials. In 11 trials, it was erroneously identified as trustworthy. However the buyer was able to realize its mistake which is indicated by the fact that its opinions were utilized less than 1% in each of the 11 trials. Lying friend was correctly identified as not trustworthy in 93 of 100 trials. It was incorrectly identified as trustworthy in 7 trials and its opinions were utilized 4% in each of those 7 trials.

Between the two methods, a buyer utilizing the second method of identifying trusted friends had a superior performance compared to a buyer using the first method of identifying trusted friends in our testing environment where majority of friends were honest. Our results show that a buyer using the second method of identifying trustworthy performs inferiorly to a buyer utilizing the first method of identifying trustworthy friends in detecting the nature of lying friends and leads us to conclude that its performance will be worse than a buyer utilizing the first method in adversarial conditions where majority of friends are either erratic or deliberately lie.

In Phase IV, we analyzed the performances of buyers with risk taking and conservative approaches. A buyer with risk tasking attitude considers a new seller as reputable initially and tends to purchase from a new seller if they are offering the lowest price among reputable sellers. A buyer with conservative attitude is cautious in its approach and explores new sellers at a rate proportional to the ratio of unexplored sellers to all the sellers who have sent bids. We compared the performances of risk taking and conservative buyers when they are acting alone and when they have friends. Metric used was: percentage of purchases from honest, high quality sellers at different levels of purchase, and as the number of sellers in the market was varied. We wanted to see how attitude affects performance of buyers when they are acting alone and when they have friends. When buyers used the second method to identify trusted friends, performance of a buyer with risk taking attitude was slightly better than the performance of a buyer with conservative attitude at lower number of purchases and at higher number of sellers. On the other hand, when buyers were acting alone or using the first method of identifying trustworthy friends, performance of buyers with conservative attitude was better than buyers with risk taking attitude at lower number of purchases, and at higher number of sellers.

In Phase V, we compared performances of all buyers (risk taking, conservative, acting alone, and with friends) at different levels of purchases, and as the number of sellers was varied. We compared all buyers on the following metric: percentage of purchases from honest, high quality sellers at different levels of purchase, and as the number of sellers in the market was varied. We wanted to see if any particular combination of buyer attitude and seller modeling yields the best performance. Our results show that when buyers are making decisions based on their own information, a buyer with conservative attitude has the best performance. When buyers are utilizing information provided by their trusted friends, a buyer with risk taking attitude and using the second method of identifying trusted friends had the best performance. Buyers with conservative and risk taking attitudes utilizing only the trusted friends' opinions fared better than buyers using trusted friends' opinions and ratings of sellers.

Chapter 5 Conclusion and Future Work

5.1 Summary

In this dissertation we presented strategies for buyers to choose sellers in decentralized, open, dynamic, uncertain and untrusted multi agent based electronic markets. We considered a marketplace where the behavior of sellers and buyers can vary, sellers and buyers can enter and leave the market any time, and sellers may be dishonest. The buyer models sellers based on its direct interactions with them, and may exchange seller information with other buyer friends in the market. There is no guarantee that when other buyers provide information, they are truthful or share similar opinions or have similar seller rating scales.

First we presented a method for a buyer to model seller reputation based on direct interactions. The buyer computes a seller's reputation based on its ability to meet its expectations of product quality and price as compared to its competitors. We show that a buyer acting alone, utilizing our model of maintaining seller reputation and buying strategy, does better than buyers acting alone employing strategies proposed previously by other researchers for frequent as well as for infrequent purchases.

Next we presented two methods for buyers in an electronic market to identify other trustworthy buyer friends, who are honest and have similar opinions regarding sellers. In both methods, the buyer models other buyers who provide seller information when it requests for it. In the first method, the buyer utilizes friends' opinions and ratings of sellers to identify trustworthy friends. The reputation of sellers provided by trustworthy friends are adjusted to account for the differences in the rating systems between the buyer and its trustworthy friends, and then combined with the buyer's own information on sellers to choose high quality, low priced sellers. In the second method, the buyer utilizes only friends' opinions of sellers to identify trustworthy friends. Ratings are assigned to sellers based on trustworthy friend's opinions and combined with the buyer's own rating on sellers to choose a high quality, low priced seller to purchase from.

We conducted experiments to show that both methods are successful in distinguishing between trustworthy friends whose opinions should be utilized in decision making and untrustworthy friends who are either dishonest, or have different opinions, and whose opinions should not be used in decision making. For the first method, our results show that honest friends with similar opinions were identified as trustworthy friends in all the 100 trials, they were among the earliest to be identified as trustworthy, and their opinions were utilized 76% to 96% of all transactions in each trial. Honest friends with slightly different opinions were also identified as trustworthy in all the trials, and their opinions were utilized 72-84% of all transactions in each trial. Honest friends with very different opinions were erroneously identified as trustworthy in 14% of our trials, and their opinions were utilized in 3% of all transactions in trials where they were identified as trustworthy. Lying friends were erroneously identified as trustworthy in 3% of 100 trials, and their opinions were utilized in less than 1% of all transactions in trials where they were identified as trustworthy.

For the second method, our results show that honest friends with similar opinions were identified as trustworthy in all the 100 trials, they were among the earliest to be identified as trustworthy, and their opinions were utilized in 99% of all transactions in each trial. Honest friends with slightly different opinions were also identified as trustworthy in all the trials, and their opinions were utilized 91% to 99% of all transactions in each trial. Honest friends with very different opinions were erroneously identified as trustworthy in 11% of 100 trials, and their opinions were utilized less than 1% of all transactions in trials where they were identified as trustworthy. Lying friends were erroneously identified as trustworthy in 7% of 100 trials, and their opinions were utilized in 4% of all transactions in trials where they were identified as trustworthy.

We showed that once trustworthy friends have been identified and their opinions have been utilized in decision making, buyers using our models of identifying friends have superior performance than a buyer acting alone at lower levels of purchases and for increasing numbers of sellers in the market. Between the two methods, a buyer utilizing the second method of identifying trusted friends had a superior performance compared to a buyer using the first method of identifying trusted friends. We also conducted experiments to determine the optimum range of the incrementing and decrementing factors for a friend's reputation.

In this dissertation we also considered risk taking and conservative attitudes of buyers. A buyer with risk tasking attitude considers a new seller as reputable initially and tends to purchase from a new seller if they are offering the lowest price among reputable sellers. A buyer with conservative attitude is cautious in its approach and explores new sellers at a rate proportional to the ratio of unexplored sellers to all the sellers who have sent bids. Our results show that, when the buyers are making decisions based on their own information, a buyer with conservative attitude has the best performance. When the buyers are utilizing information provided by their trusted friends, a buyer with risk taking attitude and utilizing only the trusted friend's opinions of sellers (second method of identifying trusted friends) has the best performance.

5.2 Future Research Directions

- Currently in our model, the buyer gives equal importance to its own impressions about a seller and information gathered from trusted friends about that seller. We would like to study the effect of varying the importance given to each component in determining the net rating of sellers.

- This dissertation has focused on strategies that buyers can use to find desirable sellers. When sellers are new to the market or to the buyers, conservative buyers may not buy immediately from them. New sellers may have to wait for a while before a buyer decides to give them a chance. This may again depend on the seller's pricing strategy. In Fire [24], researchers have suggested certified reputation. In certified reputation, an agent gathers its ratings from its interaction partners and provides them to a potential interaction partner. However, sellers who are new to the market have to still establish their reputation. In [49], Tran uses reinforcement learning to vary the seller's price to maximize the seller's profits. We would like to explore what sellers can do to in order to increase their profits, attract new buyers, and maintain their profits.

- The buyer strategies presented in this dissertation and the future works suggested so far have been based on fixed prices by sellers. We would like to extend this work by incorporating negotiation strategies on both buyer and seller sides.

- We would like to develop an interface to the agents so that human buyers and sellers can customize their agents' characteristics. For example, in the case of buyers, it could be attitude towards risk, negotiating strategy, the buyer's own knowledge of the market to be transferred to his/her agent. In the case of sellers, it could be sellers pricing strategy. In [9] Chavez et. al. conducted real life experiments with Kasbah [8]. Kasbah is an electronic market place where seller and buyer agents autonomously negotiate with each other on behalf of their human counterparts to achieve their desired goals. In [9] Chavez et. al reported the results of conducting a day's experiment with people acting as buyers or sellers and instructing their agents to buyer and sell items. We also would like to do a real life experiment with our electronic market place.

References

1. A. Abdul-Rahman and S. Hailes, “Supporting Trust in Virtual Communities”, Proc. *Hawaii International Conference on System Sciences*, Maui, Hawaii, 2000.

2. Y. Bakos, “The Emerging Role of Electronic Marketplaces on the Internet”, *Communications of the ACM*, Vol. 41(8), 35-42, August 1998.

3. S. Beldona and C. Tsatsoulis, “ An Investigation of Sharing Seller Reputation among Buyers in Agent-Based Markets”, *Third Int. Conf. on Autonomic and Autonomous Systems (ICAS)*, Athens, Greece, 2007.

4. S. Beldona and C. Tsatsoulis, “Reputation Based Buyer Strategy for Seller Selection for Both Frequent and Infrequent Purchases”, *Int. Conf. on Informatics in Control, Automation, & Robotics (ICINCO),* Angers, France, 2007.

5. J. Carbo, J. Molina, and J. Davila, “Comparing Predictions of SPORAS vs. a Fuzzy Reputation Agent System”, *Third Int .l Conf. on Fuzzy Sets and Fuzzy Systems, Interlaken*, 147–153, 2002.

6. J. Carter, E. Bitting, and A. Ghorbani, “Reputation Formalization for an Information-Sharing Multi-Agent System”, *Computational Intelligence,* Vol. 18(2), 515–534, 2002.

7. C. Castelfranchi, and R. Falcone, “Principles of Trust for MAS: Cognitive Anatomy, Social Importance, and Quantification”, Proc. *of the Intel Conf. on Multi-Agent Systems (ICMAS’98),* Paris, France, 72–79, 1998.

8. A. Chavez and P. Maes, “Kasbah: An Agent Marketplace for Buying and Selling Goods”, *First Int. Conf. on the Practical Application of Intelligent Agents and Multi-Agent Technology,* London, UK, 1996.

9. A Chavez, D. Dreilinger, R. Guttman, and P. Maes, “A Real-Life Experiment in Creating an Agent Marketplace”, *Proc. of the Second Int. Conf. on the Practical Application of Intelligent Agents and Multi-Agent Technology*, 1997.

10. C. Dellarocas, “Immunizing Online Reputation Reporting Systems Against Unfair Rating and Discriminatory Behavior”, *E-Commerce*, October 2000.

11. C. Dellarocas, “Reputation Mechanisms” *http://www.rhsmith.umd.edu/faculty/cdell/papers/elsevierchapter.pdf.* 10 Sep. 2007

12. C. Dellarocas, “Towards Incentive-compatible Reputation Management”, *Workshop on Deception, Fraud and Trust in Agent Societies, Bologna, Italy*, 26–40, 2002.

13. J.M. De Figueiredo, “Finding Sustainable Profitability in Electronic Commerce,” *Sloan Management Review*, 41-52, Summer 2000.

14. R.B Doorenbos, O. Etzioni, and D. Weld, "A Scalable Comparison-Shopping Agent for the World Wide Web", *First Int. Conf. on Autonomous Agents*, 39-48, 1997.

15. B. Esfandiari and S. Chandrasekharan, "On How Agents Make Friends: Mechanisms for Trust Acquisition", Proc. *of the Fifth Int. Conf. on Autonomous Agents Workshop on Deception, Fraud and Trust in Agent Societies*, 27-34, 2001.

16. T. Finin and B. Grosof, "Artificial Intelligence for Electronic Commerce", Proc. *of the Sixteenth National Conf. on Artificial Intelligence Workshop on Artificial Intelligence for Electronic Commerce*, 1999.

17. T. Finn, Y. Labrou, and J. Mayfield, "KQML as an Agent Communication Language," In: *Software Agents*, J. Bradshaw (Ed.), MIT Press, 1997.

18. C. Goldman, S. Kraus, and O. Shehory, "Equilibria Strategies for Selecting Sellers and Satisfying Buyers," In: *Lecture Notes in Artificial Intelligence*, M. Klusch and F. Zambonelli (Eds.), Springer, Vol. 2182, 166-177, 2001.

19. T. Grandison, and M. Sloman, "A Survey of Trust in Internet Application", *IEEE Communications Surveys*, Vol. 4(4), 2-16, 2000.

20. A. Greenwald and J. Kephart, "Shopbots and Pricebots", *Proc. of the Sixteenth Int. Joint Conf. on Artificial Intelligence*, Vol. 1, 506-511, 1999.

21. A. Greenwald, J. Kephart, and G. Tesauro, “Strategic Pricebot Dynamics”, *Proc. of the First ACM Conf. on E-Commerce*, 58-67, 1999.

22. R.H. Guttman and P. Maes, “Agent-mediated Electronic Commerce: A Survey”, *The Knowledge Engineering Review*, Vol. 13(2), 147-159, 1998.

23. D. Huynh, N.R. Jennings, and N.R. Shadbolt, “Developing an Integrated Trust and Reputation Model for Open Multi-Agent Systems”, *Proc. of 7th Int. Workshop on Trust in Agent Societies*, 65-74, New York, USA, 2004.

24. T.D. Huynh, N.R. Jennings, and N.R. Shadbolt, “An Integrated Trust and Reputation Model for Open Multi-Agent Systems”, *Autonomous Agents and Multi-Agent Systems*, Vol. 13(2), 119-154, Sep. 2006.

25. M. He, N. R. Jennings, and H. Leun. "On Agent-Mediated Electronic Commerce" *IEEE Trans on Knowledge and Data Engineering, Vol*. 15(4) 985-1003, 2003.

26. N.R. Jennings, “An Agent-based Approach for Building Complex Software Systems”, *Communications of the ACM* , Vol. 44(4), 35–41, 2001.

27. C. Jonker and J. Treur, “Formal Analysis of Models for the Dynamics of Trust Based on Experiences”, Proc*. of the Fourth Int. Conference on Autonomous Agents Workshop on Deception, Fraud and Trust in Agent Societies*, 81-94, 1999.

28. A. Jøsang, R. Ismail, and C. Boyd, "A Survey of Trust and Reputation Systems for Online Service Provision", *Decision Support System,* Vol. 43(2), 618-644, March 2007.

29. R. Jurca, and Faltings, B. (2003). "Towards Incentive-Compatible Reputation Management", In: *Trust, reputation and security: theories and practice*, Lecture Notes in AI, R. Falcone, S. Barber, L. Korba, & M. Singh, (Eds.), Springer-Verlag, Vol. 2631, 138–147, 2003.

30. B. Krulwich, "The Bargainfinder Agent: Comparision Price Shopping on the Internet", *Bots, and other Internet Beasties,* J. Williams (Ed.), 257-263, Macmillan Computer Publishing, 1996.

31. D. Kuokka and L. Harada, "Supporting Information Retrieval via Matchmaking,*" Working Notes of the AAAI Spring Symposium on Information Gathering from Heterogeneous, Distributed Environments,* 1995.

32. P. Maes, "Artificial Life Meets Entertainment: Life like Autonomous Agents," *Communications of the ACM*, Vol. 38(11), 108-114, 1995.

33. P. Maes, R.H. Guttman, and A.G. Moukas, "Agents that Buy and Sell", *Communications of the ACM*, Vol. 42(3), 81-91, 1999.

34. S. Marsh, "Formalizing Trust as a Computational Concept", *Ph.D. Thesis*, Department of Mathematics and Computer Science, University of Stirling, 1994.

35. Y. Mass, and O. Shehory, " Distributed Trust in Open Multi-Agent Systems", In: *Trust in cyber-societies,* R. Falcone, M. Singh, & Y. Tan, (Eds.), Springer-Verlag, 159–173, 2001.

36. N.H. Miller, P. Resnick, and R.J. Zeckhauser, "Eliciting Honest Feedback in Electronic Markets", KSG Working Paper Series RWP02-039, 2002.

37. S. D. Ramchurn, T. D. Huynh, and N. R. Jennings, "Trust in Multi-Agent systems", *The Knowledge Engineering Review*, Vol. 19(1), 1-25, 2004.

38. K. Regan, T. Tran, and R. Cohen, "Sharing Models of Sellers amongst Buying Agents in Electronic Marketplaces", *Decentralized Agent Based and Social Approaches to User Modelling (DASUM),* July 2005.

39. K. Regan, and R. Cohen, "A Model of Indirect Reputation Assessment for Adaptive Buying Agents in Electronic Markets", *Business Agents and the Semantic Web (BASeWEB),* May 2005.

40. S. Russell and P. Norvig, Aritificial Intelligence: A Modern Approach. Prentice Hall, 1995.

41. J. Sabater and C. Sierra " REGRET: A Reputation Model for Gregarious Societies", *Proc. of the Fifth Int. Conf. on Autonomous Agents*, Montreal, Canada, 194-195, 2001.

42. J. Sabater, "Trust and Reputation for Agent Societies", Phd Thesis, Universitat Autnoma de Barcelona, 2003.

43. J. Sabater and C. Sierra, "Review on Computational Trust and Reputation Models", *Artificial Intelligence Review*, Vol. 24(1), 33-60, 2005

44. T. Sandholm, "Distributed Rational Decision Making", In: *Multi-Agent Systems: A Modern Approach to Distributed Artificial Intelligence*, Weiss, G. (ed.), MIT Press, 1999.

45. M. Schillo, M. Rovatsos, and P. Funk, " Using Trust for Detecting Deceitful Agents in Artificial Societies", *Special Issue of the Applied Artificial Intelligence Journal on Deception, Fraud and Trust in Agent Societies*, Vol.14(8), 825–848, 2000.

46. S. Sen and N. Sajja, "Robustness of Reputation-based Trust: Boolean Case", *Proc. of the First Int. Joint Conf. on Autonomous Agents and Multiagent Systems (AAMAS-02*), Bologna, Italy, 288–293, 2002.

47. R. G. Smith, "The contract Net Protocol: High-Level Communication and Control in a Distributed Problem Solver", *Distributed Artificial Intelligence*, 357–366, 1988.

48. R. Till, "Transforming the way we do business", *In: Electronic Commerce,* T.Nash, editor, 9-12, Kogan Page, London, UK, 1998.

49. T. Tran, "Reputation-Oriented Reinforcement Learning Strategies for Economically-Motivated Agents in Electronic Market Environment", *Ph.D. Dissertation*, University of Waterloo, 2003.

50. T. Tran and R. Cohen, "Learning Algorithms for Software Agents in Uncertain and Untrusted Market Environments", Proc. *of the Eighteenth Int. Joint Conf. on Artificial Intelligence*, 2003.

51. W. T. L. Teacy, J. Patel, N.R. Jennings, and M. Luck, "Coping with Inaccurate Reputation Sources: Experimental analysis of a probabilistic trust model", *Proc of Fourth Int. Joint Conf. on Autonomous Agents and Multiagent Systems*, 997–1004, 2005 .

52. G.Weiss, Editor. Multi-Agent Systems: A Modern Approach to Distributed Artificial Intelligence, The MIT Press, 2000.

53. J.M. Vidal and E.H. Durfee, "The impact of Nested Agent Models in an Information Economy", *Second Int. Conf. on Multi-Agent Systems,* 377-384, 1996.

54. M. Wooldridge, An Introduction to MultiAgent Systems, Chichester: John Wiley and Sons, 2002.

55. M. Wooldridge and N. Jennings, "Intelligent Agents: Theory and Practice", *The Knowledge Engineering Review*, Vol. 10(2), 115-152, 1995.

56. EBay, www.ebay.com .

57. Amazon, www.amazon.com .

58. B. Yu, and M. P. Singh, “An Evidential Model of Distributed Reputation Management”, *Proc. of First Int. Joint Conf. on Autonomous Agents and Multi-Agent Systems*. Vol. 1, 294–301, 2002.

59. B. Yu and M.P. Singh, “A Social Mechanism of Reputation Management in Electronic Communities”, *Cooperative Information Agents (CIA),* Vol. 24, 154-165, Boston, USA, 2000.

60. B. Yu and M. Singh, “ Detecting Deception in Reputation Management.”, *Autonomous Agents and Multi-Agent Systems*, 73–80, 2003.

61. G. Zacharia and P. Maes, “Trust Management through Reputation Mechanisms”, *Applied Artificial Intelligence*, Vol. 14(9), 881-907, 2000.