Engendering Trust in Buying and Selling Agents by Discouraging the Reporting of Unfair Ratings

Jie Zhang
David R. Cheriton School of Computer Science
University of Waterloo
Waterloo, ON, Canada N2L 3G1
j44zhang@uwaterloo.ca

Robin Cohen
David R. Cheriton School of Computer Science
University of Waterloo
Waterloo, ON, Canada N2L 3G1
rcohen@uwaterloo.ca

Abstract

In this paper, we examine the application of electronic marketplaces, populated by buying and selling agents representing their human users, learning about potential business partners and making recommendations to their users. We propose a novel incentive mechanism to address the unfair rating problem arising when modeling the trustworthiness of selling agents relies on propagation of ratings provided by buying agents. In our mechanism, buying agents model other buyers and select the most trustworthy ones as their neighbors from which they can ask advice about sellers. In order to build reputation, sellers will model the reputation of buyers based on the number of their neighborhoods and increase values of products to satisfy reputable buyers. In consequence, our mechanism creates incentive for buyers to provide fair ratings of sellers. We also discuss how a marketplace designed in this way leads to better profit both for buyers and sellers and as such fosters trust between the agents and their human users.

Introduction

In the application area of electronic marketplaces, software agents are empowered to conduct purchasing or selling on behalf of human users. While the goals of the users can be articulated clearly, the software agents exist in an uncertain environment, where the agents serving as business partners are self-interested and therefore may engage in deception. The challenge for these software agents is to attempt to learn the behaviour of the other agents in the environment, in order to make effective decisions on behalf of their users.

Our research is aimed at investigating unfair reporting of seller reputation from one buying agent to another, when social networks of agents are created, to provide valuable information to buying agents with little experience in the marketplace. Buying agents are faced with the task of determining the trustworthiness of the agents (called advisors) that are reporting the reputability of sellers. As a result of this modeling, various neighborhoods of buying agents may emerge.

Our first insight is that if honesty is promoted in the marketplace, the problem of unfair ratings will be diminished. As a result, buying agents will share information about sellers fairly, allowing for successful sales. In addition, sellers that are trustworthy will have an accurate reporting of their worth, resulting in continued opportunities for sales. Users of these agents will therefore also be satisfied with the decisions being made on their behalf and the trust in these agents will be secured.

Our second insight is that if buying agents are to make purchasing decisions with more precise knowledge of selling agents, it will be valuable to allow information sharing within the marketplace. The possibility for unfair ratings from advisors can be mitigated by modeling these advisors in terms of both their public (view of all buying agents in the marketplace) and private (view of the buying agent, from its own past experience) reputation. It is in fact possible to carefully combine these factors and weight them according to the tolerance of the user and the confidence in his or her own knowledge of the sellers. As a result, unfair ratings can be discounted and the buying agents acting on behalf of users will continue to propose reasonable purchases for their humans. The combination of public and private knowledge is applicable both to the modeling of advisors and of sellers; this kind of framework allows for reasoning to be aligned with the preferences of the user and as such provides for a healthy relationship between the user and his or her agent, as well.

In the following sections, we sketch our proposed mechanism for addressing unfair ratings in electronic marketplaces and argue that it provides incentive for honesty. We follow this with a discussion of related work, including a contrast between our approach and those of other researchers, emphasizing how we enable trust between users and their software agents for this vital environment of electronic commerce.

E-marketplace Setting

The electronic marketplace environment we are modeling is populated with self-interested agents. Selling agents sell products to buying agents and try to maximize their profit and buying agents try to gain good products in terms of, for example, high quality and low prices. There is also a central server, which collects and maintains information about buying and selling agents, including, for example, ratings of selling agents. Through this central server, buying agents can collaborate and share ratings of selling agents. Selling agents can also make use of information about buying agents maintained by the central server, in order to distinguish them.

The buying and selling process is operated as a procurement auction where the auctioneer is a buyer and bidders are sellers. More specifically, a buying agent sends to the central server a request containing information about the product it wants to buy. The information includes the buyer's evaluation criteria for the product, which is a function of price and non-price features of the product (delivery time, for instance). In this way selling agents are able to know the buyer's values of their products. The central server forwards the request to selling agents. We assume that selling agents have registered to the central server. Sellers that are interested in selling the product to the buyer will join the procurement auction by submitting bids that describe their settings for prices of the product and values of corresponding non-price features. The auction[1] is similar to Request For Quote (RFQ) introduced by Shachat and Swarthout (Shachat & Swarthout 2003), except that RFQ is an English auction and we use a first-price sealed auction for the purpose of saving communication costs of agents. As also pointed out, an RFQ auction is equivalent to a first-price sealed auction.

The buying agent determines the winner of the auction whose product described in its bid has the highest valuation based on the buyer's evaluation criteria. It then pays the winning seller the amount, which is the price in the seller's bid. The winning seller is supposed to deliver the product to the buyer after it receives the payment. However, it may decide to alter the quality of the product actually delivered to the buying agent, or not to deliver the product at all. The buying agent finally submits a rating to the central server to report the result of the current business with the selling agent. We assume that a buyer can examine the quality of the product it purchases only after it receives the product. We also assume that there is no complete contract or legal verification to protect buying agents from dishonest sellers.

Incentive Mechanism

To formalize the proposed incentive mechanism, we consider the scenario that in an electronic marketplace a buying agent B wants to buy a product p. It sends the request to the central server. The request contains information of the buyer's evaluation criteria for a set of non-price features $\{f_1, f_2, ..., f_m\}$, as well as a set of weights $\{w_1, w_2, ..., w_m\}$ that correspond to each non-price feature. Each weight represents how much its corresponding non-price feature is worth. A higher weight for a non-price feature implies that the buyer cares more about the feature. The buyer also provides information in its evaluation criteria about the conversion from descriptive non-price feature values to numeric values (for example, 3 year warranty is converted to the numeric value of 10 on a scale of 1 to 10). We define the function $D()$ to denote such conversion. Inspired by (Boutilier, Sandholm, & Shields 2004), we also use a Quasi-linear function to represent the buyer's valuation for the product as follows:

$$V(p) = \sum_{i=1}^{m} w_i D(f_i) - P(p) \quad (1)$$

where $P(p)$ is the price of the product p.

The central server forwards the request to selling agents in the marketplace. Sellers $\overline{S}$ that are interested in selling the product to the buyer can submit their bids containing their setting for prices of the product, as well as values for non-price features. We formalize how a seller should bid for the buyer's request in the next section.

Seller Bidding for Buyer's Request

A seller S sets the price and values for the non-price features of the product p, depending on how much instant and expected future profit it can earn from selling p to the buyer B. The instant profit is the profit earned by the seller from the current transaction if it wins the auction. We define the seller's instant profit from selling the product p to the buyer B as follows:

$$U(p) = P(p) - C(p) \quad (2)$$

where $C(p)$ is the cost for the seller to produce the product p with certain values for the non-price features in its bid. The expected future profit the seller can earn depends on the reputation of the buyer. A reputable buyer in this case is one of the neighbors of many other buyers. Cooperating with reputable buyers will allow the seller to build its reputation and to be known as a trustworthy seller by many buyers in the marketplace. It will then be able to obtain more opportunities of doing business with buyers and gain more profit in the future.

To gain profit from each possible transaction, the seller may not include in its bid the true cost of producing product p with certain non-price features. Therefore, it is reasonable to assume that $P(p) > C(p)$. We define the potential gains of the buyer from the transaction as follows:

$$V'(p) = \sum_{i=1}^{m} w_i D(f_i) - C(p) \quad (3)$$

where f_i, $D()$, and w_i are defined earlier in the "Incentive Mechanism" section. We also define the distribution function for $V'(p)$ as $F(V')$.

As argued in (Shachat & Swarthout 2003), a symmetric Bayes-Nash equilibrium can be derived. The equilibrium bidding function of the seller can be derived as follows:

$$P^*(p) = C(p) + \frac{\int_{V_L - C_H}^{V'(p)} F(x)dx}{F(V')} \quad (4)$$

where V_L is the lower bound of the value for the non-price features of p and C_H is the higher bound of the cost for the seller to produce p. We assume $V_L \geq C_H$ to ensure that the value of a seller's product always exceeds its cost.

[1] Note that alternative auctions can also be deployed, such as English auction with Bidding Credits (EBC) (Shachat & Swarthout 2003). However, the study of an alternative auction is outside the scope of this paper.

By taking into account the reputation of the buyer B, the seller has the expected future profit from winning the current auction. It will reduce the instant profit and gain more chance to win the auction if the minimum expected future profit is no less than the loss of the instant profit. The bidding function of the seller in Equation 4 then should be changed to be:

$$P^*(p) = C(p) + \frac{\int_{V_L - C_H}^{V'(p)} F(x)dx}{F(V')} - V_D(R) \qquad (5)$$

where $V_D(R)$ is the valuation of discount for the buyer B with reputation $R(B)$. Comparing Equations 4 and 5, the bidding price of the seller in Equation 5 will be decreased if $V_D(R)$ is greater than 0. The buyer's valuation for the product p will then be increased, according to Equation 1. The seller will be more likely to win the auction. It is also obvious that if the bidding price is fixed, the values of the product's non-price features in the seller's bid will be increased.

Reputation of Buyer As will be discussed in the "Social Network" section, each buyer maintains a list of neighbors that it trusts the most. A seller then can model the reputation of a buyer based on the number of its neighborhoods (other buyers that include the buyer in their neighbor lists). The seller S periodically acquires neighbor list information of buyers from the central server. It then counts for each buyer the number of neighborhoods. Suppose that there are N_B other buyers considering the buyer B as one of their neighbors. The reputation of B can be calculated as follows:

$$R(B) = \begin{cases} \frac{N_B}{\theta} & \text{if } N_B < \theta; \\ 1 & \text{otherwise.} \end{cases} \qquad (6)$$

The value of θ depends on the total number of buyers in the marketplace. The buyer will be considered as reputable if $R(B)$ is no less than a threshold δ'. The buyer will be considered as disreputable if its reputation is no larger than a threshold γ' ($0 < \gamma' < \delta' < 1$).

There may exist collusion where dishonest buyers treat each other as neighbors and form a dishonest social network. We cope with this problem by allowing the seller to model the trustworthiness of a buyer. We allow the seller to check its ratings provided by the buyer. If the buyer always provides fair ratings for the seller, the buyer will be considered as a trustworthy buyer by the seller. The seller maintains a trustworthy buyer list. Based on the assumption that a trustworthy buyer's neighbors are also likely trustworthy, the seller uses the list as a basis to find other trustworthy buyers by searching the social network of buyers. From the list of all possible trustworthy buyers that the seller can find, the seller then can correctly model the reputation of a buyer.

Buyer Choosing Winning Seller

After receiving sellers' bids, the buyer B will then determine the winner of the auction. The winner of the auction is the seller whose bid includes the highest valuation of the product p that it is willing to offer, which can be formalized as follows:

$$S_{win} = arg \max_{S \in \overline{S}} V(p) \qquad (7)$$

The buyer chooses the winner of the auction among only sellers that are considered to be trustworthy. A trustworthy seller here always delivers products with values that are at least the same as what are described in its winning bids. If there are no trustworthy sellers submitting bids, the winner of the auction will be selected among the sellers that are not untrustworthy. In another words, untrustworthy sellers will be permanently barred from the buyer's auctions. Our idea of selective tendering is also supported by Kim's investigation results demonstrated in (Kim 1998). Kim states that public tendering could foster opportunism of quality reduction by bidders; in contrast, selective tendering depending on bidders' trustworthiness may avoid such opportunism.

Trustworthiness of Seller As an important component of our proposed marketplace model, the buyer models trustworthiness of a seller by using a personalized approach. It models private reputation of the selling agent based on its own ratings for the seller. If the buyer does not have enough personal experience with the seller, it will ask for its neighbors' ratings of the seller. It then can derive a public reputation of the seller from ratings provided by them. The trustworthiness of the seller will be modeled by combining the weighted private and public reputation values.

Suppose that B has the rating vector $\overline{r_{B,S}}$, which contains all the ratings provided by B for the seller S. For the purpose of simplicity, we assume that a rating for S from B is binary. For example, "1" means that the selling agent delivers the product and the valuation of the product is not less than that described in its bid, and "0" otherwise.[2] In this case, the rating of "1" will be considered as a positive rating, and "0" will be considered as a negative rating. The ratings in $\overline{r_{B,S}}$ are ordered from the most recent to the oldest according to the time when they are submitted. The ratings are then partitioned into different elemental time windows $\{T_1, T_2, ..., T_n\}$. We then count the number of positive ratings $N^B_{pos,i}$ and the number of negative ratings $N^B_{neg,i}$ in each time window T_i. The private reputation of the seller S can be estimated through the beta family of probability density functions[3] as follows:

$$R_{pri}(S) = \frac{\sum_{i=1}^{n} N^B_{pos,i} \lambda^{i-1} + 1}{\sum_{i=1}^{n} (N^B_{pos,i} + N^B_{neg,i}) \lambda^{i-1} + 2} \qquad (8)$$

where λ ($0 \leq \lambda \leq 1$)is a forgetting rate. The forgetting rate is also introduced by Jøsang and Ismail (Jøsang & Ismail 2002) to deal with possible changes of the selling agent's behavior over time because old ratings will be given less weight than more recent ones. Each user can also set this

[2]We could extend our approach to accept ratings in different ranges representing how much more or less the valuation of the product that is delivered compares with that described in the seller's bid. Accordingly, the Dirichlet family of probability density functions (Gelman *et al.* 2004) would be used to represent probability distributions of ratings.

[3]More detailed discussion of the beta function can be found in (Jøsang & Ismail 2002) and (Zhang & Cohen 2006).

rate to be high or low, according to their tolerance for less recent information.

If the buying agent B does not have enough personal experience with the seller S, it will also consider ratings provided by its neighbors. The buying agent sends a request to the central server to ask for all the ratings provided by its neighbors $\{A_1, A_2, ..., A_k\}$ for the seller S. We also partition these ratings into different elemental time windows. Suppose that the neighbor A_j provided $N_{pos,i}^{A_j}$ positive ratings and $N_{neg,i}^{A_j}$ negative ratings within the time window T_i. In the same way as estimating the private reputation, the public reputation of the seller S can be calculated as follows:

$$R_{pub}(S) = \frac{[\sum_{j=1}^{k}\sum_{i=1}^{n} N_{pos,i}^{A_j}\lambda^{i-1}Tr(A_j)] + 1}{[\sum_{j=1}^{k}\sum_{i=1}^{n}(N_{pos,i}^{A_j} + N_{neg,i}^{A_j})\lambda^{i-1}Tr(A_j)] + 2} \tag{9}$$

where $Tr(A_j)$ is the trustworthiness of the neighbor A_j. Therefore, ratings provided by more trustworthy neighbors will be given more weight.

The trustworthiness of the selling agent S is estimated by combining the weighted private and public reputation values as follows:

$$Tr(S) = wR_{pri}(S) + (1-w)R_{pub}(S) \tag{10}$$

The weight w is determined by the reliability of the estimated private reputation value as follows:

$$w = \begin{cases} \frac{N_{all}^B}{N_{min}} & \text{if } N_{all}^B < N_{min}; \\ 1 & \text{otherwise.} \end{cases} \tag{11}$$

where N_{min} represents the minimum number of ratings needed for the buyer B to be confident about the private reputation value it has of S. The calculation of N_{min} can be found in (Zhang & Cohen 2006). N_{all}^B is the total number of ratings provided by B for the seller. The seller will be considered to be trustworthy only if $Tr(S)$ is no less than a threshold δ. The seller S will be considered to be untrustworthy if its trust value is no larger than a threshold γ $(0 < \gamma < \delta < 1)$.

Social Network Similar to the social mechanism proposed in (Yu & Singh 2000), our mechanism allows each buying agent to maintain a list of neighbors from which it can trust and ask advice about sellers' trustworthiness. In order for sellers to measure reputation of buyers based on the number of their neighborhoods, each buying agent can only keep a limited number of (for example, 5) neighbors. Therefore, only the most trustworthy buyers will be kept in its neighbor list.

A buying agent models trustworthiness of another buyer (an advisor) through the personalized approach formalized in (Zhang & Cohen 2006). The buying agent first models private reputation of the advisor based on their ratings for commonly rated selling agents. The advisor will have high private reputation if they have many ratings in common. When the buying agent has limited private knowledge of the advisor, the public reputation of the advisor will also be considered. The public reputation is estimated based on all ratings for the selling agents ever rated by the advisor. The advisor will have high public reputation if its ratings are the same as the majority of the other ratings of the same sellers. Finally, the trustworthiness of the advisor will be modeled by combining the weighted private and public reputation values. These weights are determined based on the estimated reliability of the private reputation, and can be set by the user of the buying agent to reflect his confidence in his personal information.

For a new buying agent, the central server randomly assigns to it some other buying agents with high public reputation as candidates for its neighbors. The new buying agent then randomly selects some candidates as its neighbors. The neighbor list will be updated periodically. Each time, the most trustworthy candidates will be selected as neighbors. The candidate list is also updated periodically. Each time, a small portion of buyers is chosen randomly as candidates from all buyers with high public reputation values.

Examples

In this section, we use some examples to demonstrate how our mechanism works. We first provide an example to demonstrate how a buyer models trustworthiness of sellers by considering ratings of sellers provided by its neighbors, and how it selects the winning seller to do business with. We then provide another example to illustrate how a seller models reputation of buyers and specifies its bids for buyers' requests according to their reputation values.

Buyer Choosing Winning Seller

Suppose that the buyer B has two non-price features for the product p that it wants to buy. The buyer specifies a weight for each non-price feature and the information about the conversion from descriptive non-price feature values to numeric values, as presented in Table 1. To prevent it from doing business with possibly dishonest sellers, the buyer B models trustworthiness of sellers and selects trustworthy ones to do business with. Suppose that the four sellers S_1, S_2, S_3 and S_4 are all willing to sell the buyer the product p and have submitted their bids. We also suppose that the buyer B previously has not done business with any one of them. Therefore the buyer B has no ratings for these sellers. The private reputation of S_1, S_2, S_3 and S_4 can be calculated according to Equation 8 as follows:

$$R_{pri}(S_1|S_2|S_3|S_4) = \frac{0+1}{(0+0)+2} = 0.5$$

The buyer B considers ratings of sellers provided by its neighbors. We assume that the buyer B has only one neighbor, which is the buyer (advisor) A. Assume that the trust value that the buyer B has of the advisor A is 0.9. Detailed examples of how a buyer models trustworthiness of an advisor can be found in (Zhang & Cohen 2006).

The ratings of the sellers S_1, S_2, S_3 and S_4 provided by the advisor A are listed in Table 2. The symbol "T" represents a sequence of time windows, in which T_1 is the most

Table 1: Buyer B's Evaluation Criteria for p

Features	Delivery Time			Warranty		
Weights	0.4			0.6		
Descriptive values	1 week	3 days	1 day	1 year	2 years	3 years
Numerical values	3	5	10	3	5	10

recent time window. To simplify the demonstration, we assume that the advisor A provides at most one rating within each time window. Note that the advisor A does not have ratings for the seller S_2 because A has not done business with S_2. There may be various possible reasons. The seller S_2 may be considered as an untrustworthy seller by the neighbor of A. The bids submitted by S_2 may not be the highest among other sellers participating in A's auctions because S_2 has higher cost for producing products or S_2 always wants to gain more instant profit for each product.

Table 2: Ratings of Sellers provided by A

T	T_1	T_2	T_3	T_4	T_5
S_1	0	0	0	1	1
S_2	-	-	-	-	-
S_3	1	1	1	1	1
S_4	1	1	1	1	0

In this example, we set λ to be 0.9, which means that the buyer B does not have much forgetting. According to Equation 9, the public reputation of the sellers can be calculated as follows:

$$R_{pub}(S_1) = \frac{\sum_{i=4}^{5} 1 * 0.9^{i-1} * 0.9 + 1}{\sum_{i=1}^{5} 1 * 0.9^{i-1} * 0.9 + 2} = 0.39$$

$$R_{pub}(S_2) = 0.5, \quad R_{pub}(S_3) = 0.83, \quad R_{pub}(S_4) = 0.72$$

Because the buyer B has not done business with any of the sellers before, the weights of the private reputation of the sellers are all 0. The trustworthiness of the sellers can be calculated by using Equation 10 as follows:

$$Tr(S_1) = 0 * 0.5 + (1 - 0) * 0.39 = 0.39$$

$$Tr(S_2) = 0.5, \quad Tr(S_3) = 0.83, \quad Tr(S_4) = 0.72$$

We set the threshold δ to be 0.7. In this case, only the sellers S_3 and S_4 will be considered as trustworthy by buyer B.

We suppose that the sellers S_3 and S_4 may have different costs of producing the product p with certain features. The bid submitted by the seller S_3 specifies that S_3 will deliver the product with 3 year warranty in three days and the price of the product is 4. The bid submitted by the seller S_4 specifies that S_4 will deliver the product with 2 year warranty in three days and the price of the product is also 4. The values of the product p in their bids are calculated as follows:

$$V(p, S_3) = 0.4 * 5 + 0.6 * 10 - 4 = 4, \quad V(p, S_4) = 1$$

The value of the product in the bid of S_4 is lower than that of S_3. Seller S_3 will be selected as the winner. Buyer B pays S_3 the price of 4. Later on, seller S_3 delivers the product. Suppose that the seller delivers the product with 3 year warranty in three days; we say that the seller is trustworthy in this transaction. Buyer B will submit a rating of "1" to the central server.

Seller Bidding for Buyers' Requests

In this example, we illustrate how a seller S_5 models reputation of buyers and specifies its bids for buyers' requests according to their reputation values. Suppose that there are 6 buyers, $\{B_1, B_2, B_3, B_4, B_5, B_6\}$. They request the same product p with two non-price features. The weight for each non-price feature and the information about the conversion from descriptive non-price feature values to numeric values are as presented in Table 1. The seller S_5 needs to decide how to bid for each buyer's request. It models the reputation of each buyer.

Table 3: Neighbors of Buyers

Buyer	Neighbors		
B_1	B_2	B_5	B_6
B_2	B_4	B_5	B_6
B_3	B_4	B_5	B_6
B_4	B_3	B_5	B_6
B_5	B_3	B_4	B_6
B_6	B_3	B_4	B_5

Assume that each buyer is allowed to have only 3 neighbors in this example. The neighbors of each buyer are listed in Table 3. The number of neighborhoods for each buyer is:

$$N_{B_1} = 0, \quad N_{B_2} = 1, \quad N_{B_3} = 3$$

$$N_{B_4} = 4, \quad N_{B_5} = 5, \quad N_{B_6} = 5$$

If we set θ to be 6, we then calculate the reputation of each buyer according to Equation 6 as follows:

$$R(B_1) = 0, \quad R(B_2) = 0.17, \quad R(B_3) = 0.5$$

$$R(B_4) = 0.67, \quad R(B_5) = 0.83, \quad R(B_6) = 0.83$$

We set δ' to be 0.8 and γ' to be 0.3. Then, the buyers B_5 and B_6 are considered as reputable buyers, and B_1 and B_2 are disreputable buyers.

According to the reputation of each buyer, seller S_5 specifies its bid for each buyer's request. The features in each bid and profit that each buyer can gain are listed in Table 4. From this table, we can see that the reputable buyers B_5 and B_6 are able to gain the largest profit and the disreputable buyers B_1 and B_2 can gain the smallest profit.

Table 4: Profit Gained by Different Buyers

Buyers	Features of Product			Profit
	Warranty	Delivery Time	Price	
B_1, B_2	1 year	1 week	5	-2
B_3, B_4	2 years	3 days	4	1
B_5, B_6	3 years	1 day	3	7

Experimental Results

We carry out experiments to examine each expectation of our mechanism. We also measure profit gained by different buyers and sellers. The expectation is that reputable buyers and sellers that are considered as trustworthy by many buyers will be able to gain more profit.

We simulate a marketplace operating with our mechanism in the period of 20 days. The marketplace involves 100 buyers. These buyers have different numbers of requests. Every 10 of them has requests in different percentages (from 10% to 100%) of the time period (20 days). We assume that there is only one product in each request and each buyer has a maximum of one request each day. For the purpose of simplicity, we also assume that the products requested by buyers have the same non-price features. After they finish business with sellers, buyers rate sellers. 50 buyers provide unfair ratings. Every 10 of them provides different percentages (from 10% to 50%) of unfair ratings. Initially, we randomly assign 5 buyers to each buyer as its neighbors.

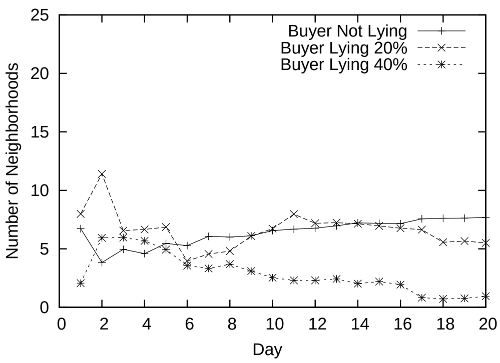

Figure 1: Reputation of Different Buyers

There are 10 sellers in total in the marketplace. Each 2 sellers acts dishonestly in different percentages (0%, 25%, 50%, 75% and 100%) of their business with buyers. One half of the sellers model reputation of buyers and adjust prices of products according to buyers' reputation. Another 5 sellers do not model reputation of buyers. They offer the same price for products requested by buyers. We assume that all sellers have the same cost for producing the products because all products have the same non-price features.

We first measure the reputation of buyers that provide different numbers of unfair ratings. The results are shown in Figure 1. In our experiments, the reputation of a buyer is represented by the number of the buyer's neighborhoods. From this figure, we can see that the buyers providing the

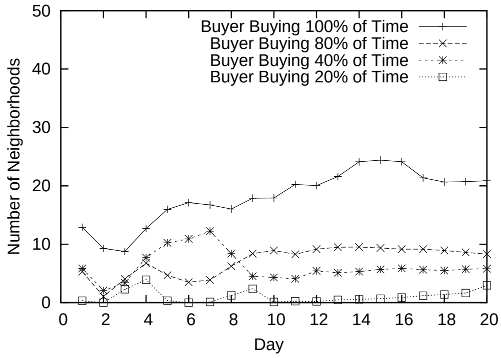

Figure 2: Reputation of Different Buyers

smaller number of unfair ratings will have the larger reputation values. Due to the randomness of the initial setting for our experiments, buyers providing more unfair ratings may have larger reputation values at the beginning. But their reputation will continuously decrease after approximately 10 days, as can be seen from Figure 1. After approximately 14 days when our marketplace converges, the buyers providing more unfair ratings will have smaller reputation values. We also measure reputation of buyers that have different numbers of requests. Results are shown in Figure 2. Buyers having more requests to purchase products will have larger reputation values. Similarly, reputation values of buyers change stochastically at the beginning. But when the marketplace converges, the buyers having fewer requests will have the smaller reputation values.

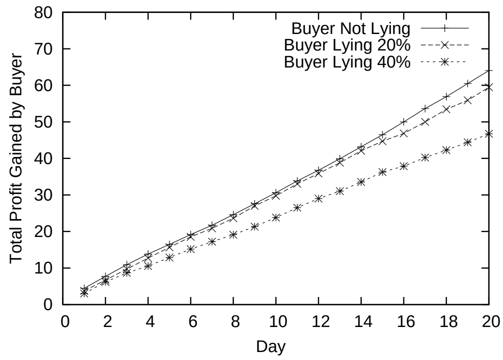

Figure 3: Profit Gained by Different Buyers

After each day, we measure total profit gained by buyers that provide different numbers of unfair ratings. The profit gained by a buyer from buying a product is calculated using Equation 1. From Figure 3, we can see that buyers providing fewer unfair ratings will gain more total profit. It is better off for buyers to provide a greater number of truthful ratings. Note that the profit difference for different types of buyers is fairly small. This is because buyers do not have many requests (at most 20). We do not measure total profit gained by buyers having different numbers of requests, because the more requests buyers have, the more profit they will be able to gain.

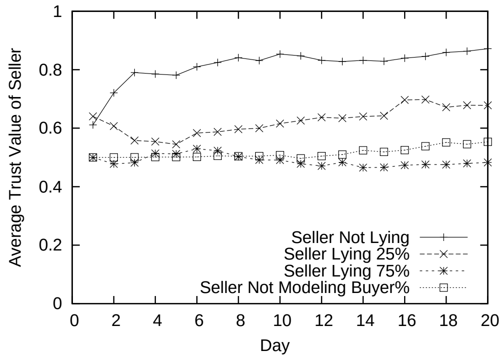

Figure 4: Average Trust Value of Different Sellers

We compare average trust values of different sellers. The average trust value of a seller is calculated as the sum of a trust value each buyer has of the seller divided by the total number of buyers. As shown in Figure 4, results indicate that sellers being dishonest more often will have smaller average trust values. The sellers that do not model reputation of buyers will also have smaller average trust values. From Figure 4, we can see that their average trust values are nearly 0.5. It is because that they do not have much chance to do business with buyers and will not have many ratings. A seller without any ratings will have trust value of 0.5 (for example, the seller S_2 in the "Examples" section). Similarly, the sellers being dishonest in 75% of their business also will not have much chance to do business with buyers and will have a trust value of nearly 0.5.

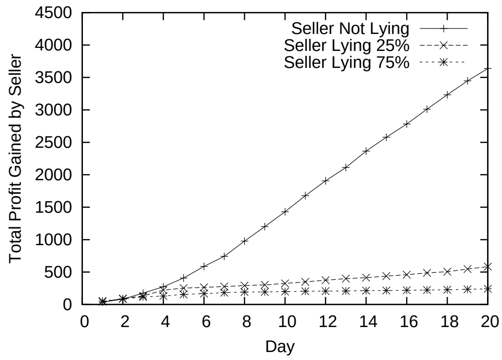

Figure 5: Total Profit Gained by Different Sellers

We also compare total profit gained by different sellers. Results are shown in Figures 5 and 6. From Figure 5, we can see that sellers being honest more often will gain more profit. Therefore, it is better off for sellers to be honest. Note that the profit difference between the honest sellers and the sellers lying 25% is much larger than that between the sellers lying 25% and the sellers lying 75%. It is because we set the threshold δ to be very high (0.8). The sellers lying 25% will not be considered as trustworthy sellers and therefore will have only a small chance to be selected as business partners. Results in Figure 6 indicate that sellers are better off to model reputation of buyers and adjust prices of products according to buyers' reputation, in order to gain more profit.

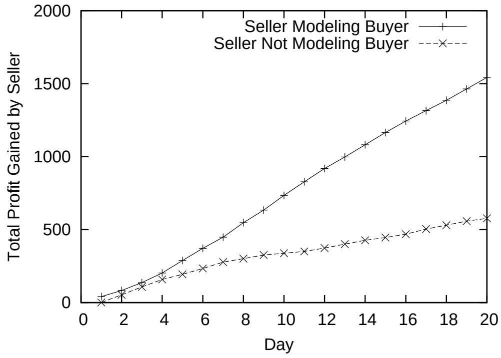

Figure 6: Total Profit Gained by Different Sellers

Discussion and Related Work

The algorithms for buying and selling agents described in this paper are ones where agents learn in a dynamic environment, in order to make valuable recommendations to their human users. With incentives for honesty in place, the chance of poor decisions being made by these agents decreases (as evidenced by the experimental results that buying and selling agents using our mechanism will increase their profit). As such, a bond of trust is built between each agent and its human user. Honesty is promoted in our mechanism as follows. Buyers report their evaluation of sellers to a central server and from here to other buying agents. This helps to keep sellers honest. Sellers can examine how they have been rated by the buyer. This helps to keep buyers honest. In addition, buyers will be rewarded if they are considered trustworthy by other buyers since sellers can increase the value of their products for buyers who reside in many neighborhoods. This is an additional incentive for honesty in reporting from buyers.

There are other approaches for promoting honesty in electronic marketplaces. Two such methods are side payments (Jurca & Faltings 2003; Miller, Resnick, & Zeckhauser 2005) and credibility mechanisms (Papaioannou & Stamoulis 2005; Jurca & Faltings 2004). Side payment mechanisms offer side payment to buyer agents that fairly rate results of business with sellers. In these mechanisms, providing fair ratings for business results is a Nash equilibrium. Credibility mechanisms measure agents' credibility. The credibility of two agents (a buyer and a seller, for example) in their business will be decreased if their ratings about the business result are different. Buyers will provide fair ratings in order to keep up their credibility. In contrast, our approach allows buying and selling agents to learn about each other, in order to make effective decisions on behalf of their human users about which sellers to buy products from and how to sell products to particular buyers. In so doing, honest buying and selling agents will gain better profit for their users, and the trust between agents and their users will be fostered.

Our buying and selling agents also have important opportunities to personalize their decision making according to the preferences of their human users. Buying agents reason with a weighting factor, w, to determine the contribution of private and public information to the overall calculation of seller trustworthiness. A user who is confident in its personal information can elect to weight that more heavily; this is also the case for a user who does not want to rely on the reports of other users (e.g. one who is very particular in his likes and needs). It would be useful to initiate a dialogue between the agent and user to set the value of w and to refine it, as the buyer learns more about the marketplace, after examining the goods he or she has purchased. A similar kind of reasoning would take place for the threshold γ used to express tolerance of untrustworthy behaviour from sellers. Selling agents must reason about a threshold as well, when classifying the buyers of the marketplace.

Future Work

One interesting topic for future research that focuses on the interaction between the users and their agents is how best to capture the users' preferences, in order to set the user-specific factors in our equations. One can imagine employing a framework for building up user modeling values as discussed in (Fleming 2004), using a combination of general user models, stereotypes and specific user modeling features, in order to represent (and adjust over time) the preferences of the users with respect to tolerance for trusting others in the marketplace.

In future work, we will also investigate the possibility of not relying on the central server, adopting instead a distributed incentive mechanism. Different issues may arise in the distributed case. A buyer may untruthfully rate a seller a large number of times, which is referred to as "flooding" the system (Dellarocas 2000). One way to cope with this problem is to develop a system to certify ratings provided by buyers. Another issue is the fact that it is costly for a buyer to obtain public knowledge of other buyers without the central server. In this case, we may allow each buyer to keep a large set of other buyers as candidates for its neighbors. Public reputation of other buyers can be measured based on ratings provided by these candidates.

Another topic for future work is to examine marketplaces where the identity of buyers is shielded from the seller, to prevent sellers from trying to cheat less reputable buyers (that do not have much impact on the seller's reputation). For instance, the seller could submit bids for certain classes of buyers to the central server and indicate its value for the reputation of each buyer. The central server could then deliver the appropriate bid to the buyer trying to purchase from this seller and keep the buyer's identity protected.

Our mechanism allows sellers to model reputation of a buyer based only on the number of other buyers including the buyer in their neighbor lists. In future work, we will consider a more comprehensive approach for modeling buyers' reputation. The reputation of buyers that include the buyer in their neighbor lists could also be taken into account. How best to form neighborhoods in the marketplace is another open question for research.

References

Boutilier, C.; Sandholm, T.; and Shields, R. 2004. Eliciting bid taker non-price preference in (combinatorial) auctions. In *Proceedings of the Nineteenth National Conference on Artificial Intelligence (AAAI-04)*.

Dellarocas, C. 2000. Immunizing online reputation reporting systems against unfair ratings and discriminatory behavior. In *Proceedings of the 2nd ACM Conference on Electronic Commerce*.

Fleming, M. 2004. The use of increasingly specific user models in the design of mixed-initiative systems. In *Proceedings of the Seventeenth Canadian Conference on Artificial Intelligence*.

Gelman, A.; Carlin, J. B.; Stern, H. S.; and Rubin, D. B. 2004. *Bayesian Data Analysis*. Chapman & Hall/CRC, Boca Raton.

Jøsang, A., and Ismail, R. 2002. The beta reputation system. In *Proceedings of the 15th Bled Electronic Commerce Conference*.

Jurca, R., and Faltings, B. 2003. An incentive compatible reputation mechanism. In *Proceedings of the IEEE Conference on E-Commerce, Newport Beach, CA, USA*.

Jurca, R., and Faltings, B. 2004. Confess: An incentive compatible reputation mechanism for the online hotel booking industry. In *Proceedings of the IEEE Conference on E-Commerce, San Diego, CA, USA*.

Kim, I.-G. 1998. A model of selective tendering: Does bidding competition deter opportunism by contractors? *The Quarterly Review of Economics and Finance* 38(4):907–925.

Miller, N.; Resnick, P.; and Zeckhauser, R. 2005. Eliciting informative feedback: The peer-prediction method. *Management Science* 51(9):1359–1373.

Papaioannou, T. G., and Stamoulis, G. D. 2005. An incentives' mechanism promoting truthful feedback in peer-to-peer systems. In *Proceedings of IEEE/ACM CCGRID 2005 (Workshop on Global P2P Computing)*.

Shachat, J., and Swarthout, J. T. 2003. Procurement auctions for differentiated goods. *Working Paper, IBM Research Labs*.

Yu, B., and Singh, M. P. 2000. A social mechanism of reputation management in electronic communities. In *Proceedings of the 4th International Workshop on Cooperative Information Agents*, 154–165.

Zhang, J., and Cohen, R. 2006. Trusting advice from other buyers in e-marketplaces: The problem of unfair ratings. In *Proceedings of the Eighth International Conference on Electronic Commerce (ICEC'06)*.