Design of a Mechanism for Promoting Honesty in E-Marketplaces

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 explore the use of the web as an environment for electronic commerce. In particular, we develop a novel mechanism that creates incentives for honesty in electronic marketplaces where human users are represented by buying and selling agents. In our mechanism, buyers model other buyers and select the most trustworthy ones as their neighbors from which they can ask advice about sellers. In addition, however, sellers model the reputation of buyers. Reputable buyers provide fair ratings of sellers, and are likely to be neighbors of many other buyers. Sellers will provide more attractive products to reputable buyers, in order to build their reputation. We discuss how a marketplace operating with our mechanism leads to better profit both for honest buyers and sellers. With honesty encouraged, our work promotes the acceptance of web-based agent-oriented e-commerce by human users.

Introduction

In electronic marketplaces that lack complete contracts and legal verification, a buying agent often relies on self-enforcing contracts where it can selectively choose business partners (selling agents) depending on their trustworthiness. A modeling of the trustworthiness of a seller can be based on the buyer's past personal experience with the seller. However, for a new buyer or a buyer without any personal experience with the seller, evaluation of the seller's trustworthiness is often determined by examining the ratings for the seller from other buyers. The problem of unfair ratings may then arise. Buyers may provide unfairly high ratings to promote the seller. This is referred to as "ballot stuffing" (Dellarocas 2000). Buyers may also provide unfairly low ratings, in order to cooperate with other sellers to drive a seller out of the marketplace. This is referred to as "bad-mouthing". A second problem is that rating submission is voluntary. Buyers do not have direct incentives to provide ratings because, for example, providing reputation ratings of sellers requires some effort (Jøsang, Ismail, & Boyd 2005; Miller, Resnick, & Zeckhauser 2005). Providing fair ratings for a trustworthy seller may also decrease the chance of doing business with the seller because of competition from other buyers.

To address these two problems, we propose a novel incentive mechanism to elicit fair ratings of sellers in electronic marketplaces. In our mechanism, buyers are encouraged to be truthful in order to gain more profitable transactions. This idea is supported by Gintis et al. (Gintis, Smith, & Bowles 2001). They argue that altruism in one context signals "quality" that is rewarded by increased opportunities in other contexts. Specifically, if the system is such that the provision of truthful reputation feedback makes agents more likely to choose to undertake transactions with the reporting agent then the reporting agent would benefit for its feedback through a greater number of profitable transactions. In our mechanism, the reputation of buyers is modeled by sellers. A buyer is considered reputable if it is well respected in the social network - i.e. it is a neighbor of many other buyers. Sellers increase quality and decrease prices of products to satisfy reputable buyers, in order to build their own reputation. Our mechanism, therefore, creates incentives for buyers to provide fair ratings of sellers. Agents are designed to aid human users in reasoning about business partners. With honesty encouraged, our work therefore assists in encouraging human users to participate in agent-oriented electronic commerce on the web.

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. We also assume a central server that accepts requests from buyers. The buyer's request contains information of its evaluation criteria for a set of non-price features $\{f_1, f_2, ..., f_m\}$ of the product, as well as a set of weights $\{w_1, w_2, ..., w_m\}$ that correspond to the features. The buyer also provides a function $D()$ to convert 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).[1] We use a quasi-linear function to represent the buyer's valuation for p as follows:

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

[1] In this paper, we focus on non-price features that are still objective - e.g. delivery time. Handling subjective features is left for future work.

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

The central server forwards the buyer's request to sellers in the marketplace. We assume that the buying and selling process is operated as a procurement auction.[2] Sellers $\overline{S}$ that are interested in selling the product to B can submit 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 $S \in \overline{S}$ sets the price and values for the non-price features of p, depending on how much instant profit it can earn from selling p to B. The instant profit is the profit earned by S from the current transaction if it wins the auction. We define the instant profit as follows:

$$U(p) = P(p) - C(p) \tag{2}$$

where $C(p)$ is the cost for S to produce p.

To gain profit from each possible transaction, the seller may not include in its bid the true cost of producing p with certain features. The best potential gain the seller can offer the buyer from the transaction is as follows:

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

We define the distribution function for $V'(p)$ as $F(V')$. A symmetric Bayes-Nash equilibrium can be derived (Shachat & Swarthout 2003). The equilibrium bidding function can be derived as follows:

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

where V_L is the lower bound of the value for the non-price features of p and C_H ($V_L \geq C_H$) is the higher bound of the cost for the seller to produce p.

By taking into account the reputation of the buyer, the seller has the expected future profit from winning the current auction. 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. The seller will reduce the instant profit (decrease the price or increase the quality of p) and gain more chance to win the auction if the 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:

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

where $V_D(R) \geq 0$ is the valuation of discount for the buyer with reputation $R(B)$.

Our mechanism allows the central server to maintain a fixed number of neighbors for each buyer: a list of the most trustworthy other buyers, used to provide advice about sellers. The central server models the trust value a buyer has of another buyer (an advisor) through the personalized approach formalized in (Zhang & Cohen 2006). It first models private reputation of the advisor based on their ratings for commonly rated sellers. When the buyer has limited private knowledge of the advisor, the public reputation of the advisor will also be considered, based on all ratings for the sellers ever rated by the advisor. Finally, the trustworthiness of the advisor will be modeled by combining the weighted private and public reputation values.

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:

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

The value of θ depends on the total number of buyers in the marketplace.[3] 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 ($0 < \gamma < \delta < 1$).

Buyer Choosing Winning Seller

After receiving sellers' bids, the buyer B will then determine the winner of the auction. The winner is the seller whose bid includes the highest valuation of the product p that it is willing to offer. The buyer chooses the winner of the auction among only sellers that are considered to be trustworthy.

The buyer models trustworthiness of a seller by also using a personalized approach. Suppose that B has provided some ratings for the seller S. We assume that a rating is binary. For example, "1" (a positive rating) means that the valuation of the delivered product is not less than that described in the seller's bid, and "0" (a negative rating) otherwise. The ratings are ordered from the most recent to the oldest according to the time when they are submitted, and are partitioned into different elemental time windows $\{T_1, T_2, ..., T_n\}$. We then count the number of positive ratings $N_{ps,i}^B$ and negative ratings $N_{ng,i}^B$ in each time window T_i. The private reputation of S can be estimated through the beta family of probability density functions (Jøsang, Ismail, & Boyd 2005) as follows:

$$R_{pri}(S) = \frac{\sum_{i=1}^{n} N_{ps,i}^B \lambda^{i-1} + 1}{\sum_{i=1}^{n} (N_{ps,i}^B + N_{ng,i}^B)\lambda^{i-1} + 2} \tag{7}$$

where λ ($0 \leq \lambda \leq 1$) is a forgetting rate for assigning less weight to old ratings, to deal with possible changes of the seller's behavior over time.

---

[2] Note that our mechanism is generally applicable to marketplaces where sellers may alter quality and prices of their products to satisfy honest buyers.

[3] For the examples in this paper, we equate θ with number of buyers. Developing more sophisticated measurements of θ is left for future work.

If the buyer B does not have enough personal experience with the seller S, it will send a request to the central server to ask for all the ratings of S provided by its neighbors $\{A_1, A_2, ..., A_k\}$. Suppose that the neighbor A_j has provided $N_{ps,i}^{A_j}$ positive ratings and $N_{ng,i}^{A_j}$ negative ratings of S 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:

$$R_{pub}(S) = \frac{\sum_{j=1}^{k}\sum_{i=1}^{n} N_{ps,i}^{A_j} \lambda^{i-1} Tr(A_j) + 1}{\sum_{j=1}^{k}\sum_{i=1}^{n} (N_{ps,i}^{A_j} + N_{ng,i}^{A_j}) \lambda^{i-1} Tr(A_j) + 2} \tag{8}$$

where $Tr(A_j)$ is the trustworthiness of A_j.

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

$$Tr(S) = wR_{pri}(S) + (1-w)R_{pub}(S) \tag{9}$$

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

$$w = \begin{cases} \frac{N_{all}^{B}}{N_{min}} & \text{if } N_{all}^{B} < N_{min}; \\ 1 & \text{otherwise.} \end{cases} \tag{10}$$

where N_{all}^{B} is the total number of ratings provided by B for S, and N_{min} represents the minimum number of ratings needed for B to be confident about the estimated private reputation value. The calculation of N_{min} can be done in a similar way as that proposed in (Zhang & Cohen 2006). 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' $(0 < \gamma' < \delta' < 1)$.

Once a buyer engages in commerce with a seller, the buyer submits its rating of the seller to the central server. This information is used by the central server to form the neighborhoods of advisors for other buyers. The information may also be viewed by the seller, in order to determine the reputability of the buyer. The rating provided by the buyer is a binary value and is a reflection of whether the buyer believes that the seller delivered fairly on its stated promise for the good.

Examples

In this section, we use some examples to demonstrate how our mechanism works.

Buyer Choosing Winning Seller

We first provide an example to demonstrate how a buyer models trustworthiness of sellers by considering ratings of them provided by its neighbors, and how it selects the winning seller to do business with. Suppose that a 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

Table 1: B's Evaluation Criteria for p

Non-price Features	Delivery Time (day)			Warranty (year)		
Weights	0.4			0.6		
Descriptive Value	7	3	1	1	2	3
Numerical Value	3	5	10	3	5	10

the information about the conversion from descriptive non-price feature values to numeric values, as presented in Table 1. To prevent from doing business with possibly dishonest sellers, 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 have submitted their bids. We also suppose that B previously has not done business with any one of the sellers. Therefore B has no ratings for these sellers. The private reputation of the sellers can be calculated according to Equation 7 as follows:

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

The buyer B then considers ratings of the sellers provided by its neighbors. We assume that B has only one neighbor, which is the buyer (advisor) A. Assume that the trust value that the buyer B has of A is 0.9. The ratings of the sellers 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 recent time window. To simplify the demonstration, we assume that A provides at most one rating within each time window. Note that the advisor A does not have ratings for S_2 because A has not done business with S_2.

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 8, 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 9 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 sellers by B.

We suppose that the sellers S_3 and S_4 may have different costs for producing the product p with certain features. The bid submitted by S_3 specifies that it will deliver the product with 3 year warranty in three days and the price of the product is 4. The bid submitted by S_4 specifies that it will deliver the product with 2 year warranty in three days and the price of the product is also 4. The values of 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 p 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 S_3 delivers the product as promised in its bid; we say that S_3 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, $\{B_1, B_2, ..., B_6\}$. They request the same product p with features and associated weights 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. We count the number of neighborhoods for each buyer as follows:

$$N_{B_1} = 0, \quad N_{B_2} = 1, \quad N_{B_3} = 3$$
$$N_{B_4} = 4, \quad N_{B_5} = 5, \quad N_{B_6} = 5$$

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 buyers, S_5 specifies its bids for their request. The features of p in each bid and the 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	7 days	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 simulate a marketplace operating with our mechanism in the period of 20 days. The marketplace involves 100 buyers. Every 10 buyers has a different number (from 2 to 20) of requests. Each buyer will submit a rating for each of its transaction with a seller. Therefore, buyers having a larger number of requests will provide a larger number of ratings. 50 buyers provide unfair ratings. Every 10 of them provides different percentages (from 10% to 50%) of unfair ratings. We assume that there is only one product in each request and each buyer has a maximum of one request each day. We also assume that the products requested by buyers have the same features. Initially, we randomly assign 5 other buyers to each buyer as its neighbors; we update this list each day to retain the 5 most trustworthy other buyers as neighbors. There are also 10 sellers in the marketplace. Each 2 sellers acts dishonestly in different percentages (0%, 25%, 50%, 75% and 100%) of their business with buyers. Half of them 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. We assume that all sellers have the same cost for producing the products.

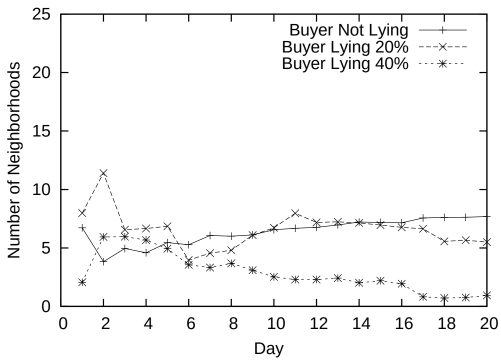

Figure 1: Reputation of Different Buyers

We first measure reputation (represented by the number of neighborhoods) of buyers that provide different numbers of unfair ratings. From Figure 1, we can see that the buyers providing the 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 11 days. 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. As shown in Figure 2, buyers having more requests (providing more ratings) will have larger reputation values.

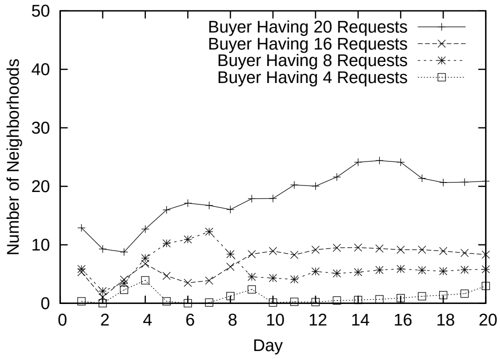

Figure 2: Reputation of Different Buyers

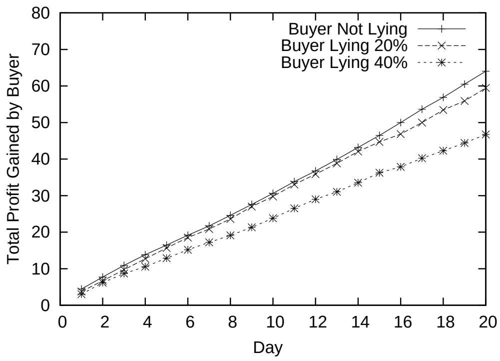

Figure 3: Total Profit Gained by Different Buyers

After each day, we measure total profit gained by different buyers. The profit gained by a buyer from buying a product is calculated using Equation 1. Shown in Figure 3, buyers providing fewer unfair ratings will gain more total profit. It is better off for buyers to provide more fair ratings. Note that the profit difference of different types of buyers is fairly small. It 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.

We compare average trust values of different sellers, which are calculated as the sum of trust values each buyer has of the sellers divided by the total number of buyers. As shown in Figure 4, sellers being dishonest more often will have smaller average trust. The sellers that do not model reputation of buyers and adjust their prices of products according to buyers' reputation will also have smaller average trust values, which are nearly 0.5. This is because they do not have much chance to do business with buyers and will not have many ratings. Similarly, the sellers being dishonest in 75% of their business also do not have much chance to do business with buyers and have trust values of nearly 0.5.

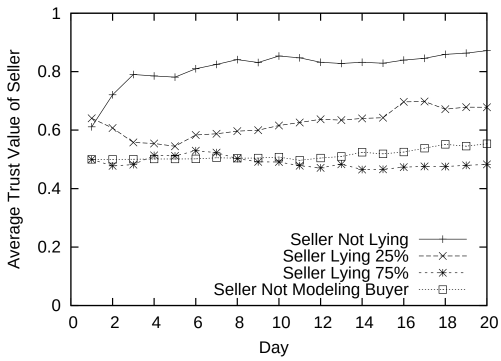

Figure 4: Average Trust Value of Different Sellers

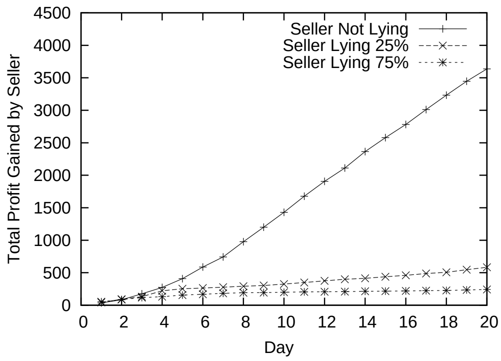

Figure 5: Total Profit Gained by Different Sellers

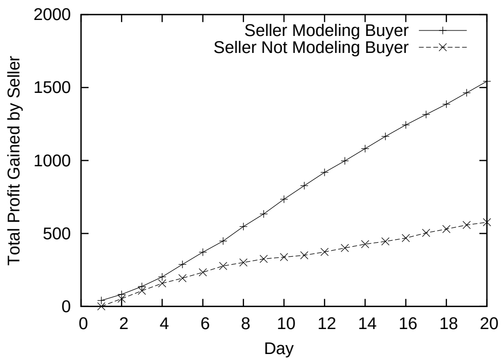

Figure 6: Total Profit Gained by Different Sellers

We also compare total profit gained by different sellers. Shown in Figure 5, 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%. This is because we set the threshold δ' to be very high (0.8). The sellers lying 25% will not be considered as trustworthy sellers, therefore will have 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.

Related Work

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 buying agents that fairly rate results of business with sellers. In these mechanisms, providing fair ratings for business results is a Nash equilibrium. Side payment mechanisms assume that buyers act independently, and therefore may have difficulty with the situation where buyers collude in giving unfair ratings. In contrast, our mechanism can allow sellers to view the ratings provided by buyers and can in this way detect dishonesty. Since sellers also only reward reputable buyers, buyers that collude in providing dishonest ratings will not profit. In addition, honest buyers will not be adversely affected by collusion in the marketplace; with our personalized approach for modeling the trustworthiness of advisors, each buyer can rely on private knowledge to detect dishonest buyers and will have their neighborhood of advisors limited to those that are determined to be trustworthy.

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. However, credibility mechanisms cannot deal with the situation where buyers and sellers collude to increase each other's credibility. Because our mechanism allows a central server to maintain a list of trustworthy neighbors for each buyer, a buyer can make an informed decision about which sellers to do business with. If a buyer were to accept the advice of another agent that is colluding with a seller and then be disappointed with the purchase, the advisor would be considered untrustworthy and would not impact any future decisions.

Conclusions and Future Work

In this paper, we propose a novel incentive mechanism to elicit fair ratings of selling agents in electronic marketplaces. In our mechanism, a buying agent maintains a neighbor list of other buyers that are considered to be trustworthy. We allow sellers to see how they have been rated by buyers and to model the reputation of buyers based on the social network. Reputable buyers are likely to be neighbors of many other buyers. Sellers then provide more attractive products to buyers that are determined to be reputable. Hence, buyers are better off providing truthful feedback and becoming neighbors of as many other buyers as possible. Sellers are also kept honest, because buyers are modeling the fairness of ratings provided by other agents, when forming their neighbor lists of other buyers. Sellers are motivated to provide quality service to reputable buyers, in order to progressively build their reputation in the social network. The above expectations are upheld in our model and shown by our experiments.

For future work, it would be useful to explore some extensions to our proposed mechanism. For example, we could develop a more detailed method for forming the neighborhoods of buyers in the marketplace. We could also introduce a more comprehensive approach for modeling a buyer's reputation based on the social network topology, by, for example, taking into account the reputation of the buyers that have this buyer included in their neighborhood list. We also plan to study the area of reputation in repeated procurement auctions to formalize sellers' future profit, in order for sellers to determine how much discount should be offered to each buyer.

We will also develop more extensive experimentation to continue to validate our model. We are particularly interested in empirically demonstrating how our framework is able to handle marketplaces where strategic agents collude with each other. This problem has been acknowledged as an important consideration by several researchers in the field (e.g (Jurca & Faltings 2003)) and if addressed in a convincing manner would continue to provide valuable encouragement for the use of web-based agents for electronic commerce.

References

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

Gintis, H.; Smith, E. A.; and Bowles, S. 2001. Costly signaling and cooperation. *Journal of Theoretical Biology* 213:103–19.

Jøsang, A.; Ismail, R.; and Boyd, C. 2005. A survey of trust and reputation systems for online service provision. *(to appear). Decision Support Systems*.

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*.

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*.

Zhang, J., and Cohen, R. 2006. A personalized approach to address unfair ratings in multiagent reputation systems. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'06) Workshop on Trust in Agent Societies*.