Towards More Effective E-Marketplaces: A Novel Incentive Mechanism

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 the context of electronic commerce, the problem of unfair ratings arises when modeling the trustworthiness of a selling agent relies (partially) on propagation of ratings provided by buying agents that have personal experience with the seller. Extreme diversity of open and dynamic electronic marketplaces causes difficulties in handling unfair ratings in trust management systems. To ease this problem, we propose a novel incentive mechanism for eliciting fair ratings of sellers from buyers. In our mechanism, buyers model other buyers, using an approach that combines both private and public reputation values. In addition, however, sellers model the reputation of buyers. Reputable buyers provide fair ratings of sellers, and are likely considered trustworthy by many other buyers. In marketplaces operating with our mechanism, sellers will offer more attractive products to satisfy reputable buyers, in order to build their reputation. In consequence, our mechanism creates incentives for buyers to provide fair ratings of sellers, leading to more effective e-marketplaces where honest buyers and sellers can gain more profit.

Introduction

In multiagent systems for electronic commerce, trust plays an important role. It provides a form of social control and allows agents in e-marketplaces to reason about reliability, capability and honesty of others, in order to choose the best business partners. Researchers have been working at designing frameworks to model the trust and reputation of agents. A modeling of the trustworthiness of a selling agent can be based on a buying agent'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 (Sabater & Sierra 2005). 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".

Besides the problem of unfair ratings, rating submission is voluntary in most trust management systems. 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, researchers have been working on developing incentive mechanisms. The aim is to encourage honesty in the reporting from buyers, in order to diminish concerns about unfair ratings. Two types of mechanisms have been developed, side payment mechanisms (Jurca & Faltings 2003; Miller, Resnick, & Zeckhauser 2005), and credibility mechanisms (Papaioannou & Stamoulis 2005; Jurca & Faltings 2004). Side payment mechanisms offer side payment to buyers 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 participants (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.

We first develop a model (a personalized approach) that addresses unfair ratings but with more flexibility for buyers to weight their value in private and public reputation ratings of other buyers (advisors). Our aim is to develop improved methods for modeling trustworthiness of advisors by tracking ratings provided according to their related time windows. In so doing, our approach is able to avoid the situation where advisors may untruthfully rate sellers for a large number of times (known as "flooding") and deal with changes of agents' behavior. Our method is able to cope with large numbers of unfair ratings.

Equipped with the richer method for modeling trustworthiness of advisors in terms of private and public reputation, we then propose a novel incentive mechanism. Our mechanism does not rely on side payment. Instead, buyers are encouraged to be truthful in order to gain more numbers of profitable transactions. This idea is supported by the work in the field of evolutionary game theory, such as the work of Gintis et al. (Gintis, Smith, & Bowles 2001). They argue that an agent's altruism in one context signals "quality" of the agent that will benefit from increased opportunities in other wider 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.

Our personalized approach provides the promising first step for our work. It allows buyers to effectively model the trustworthiness of other buyers. We then use this approach to create a social network of buyers. Each buyer in the society retains a neighborhood of the most trustworthy buyers, as advisors. In our mechanism, we also allow sellers to explicitly model the reputability of buyers, based on the number of neighborhoods they belong to in the society. Buyers that always provide fair ratings of sellers are likely to be neighbors of many other buyers and can be considered reputable. This is also supported by Gintis et al. (Gintis, Smith, & Bowles 2001) through the model of a multi-player game. They argue that agents reporting honestly provide benefit to others and will further be preferred by others as allies. These agents will be able to attract a larger number of audience to witness their feedback (also known as increasing "broadcast efficiency"). Sellers in our system will increase quality and decrease prices of products to satisfy reputable buyers. This therefore creates an incentive for buyers to provide fair ratings of sellers.

The rest of the paper is organized as follows. We first introduce the setting of the electronic marketplace in which our incentive mechanism can operate. We then formalize our mechanism and demonstrate some examples. We also describe our simulations and experimental results. After that, we introduce some related work and contrast approaches of other researchers with our work. Finally, we present conclusions and future work.

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 buyers 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 buyers and sellers, including, for example, ratings of sellers. Through this central server, buyers can collaborate and share ratings of sellers. Sellers can also make use of information about buyers maintained by the central server, in order to distinguish them.

The buying and selling process is operated as a procurement (reverse) auction where the auctioneer is a buyer and bidders are sellers.[1] More specifically, a buyer 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 sellers are able to know the buyer's values of their products. The central server forwards the request to sellers. We assume that sellers 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[2] 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 buyer determines the winner of the auction whose product described in its bid has the highest valuation based on the buyer's evaluation criteria. The buyer 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 buyer, or not to deliver the product at all. The buyer finally submits a rating to the central server to report the result of the current business with the seller. 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 buyers from dishonest sellers.

[1] We use this special setting for demonstrating our proposed approach. However, our incentive mechanism is generally applicable to marketplaces where sellers may alter quality and prices of their products to satisfy honest buyers.

Incentive Mechanism

To formalize the proposed incentive mechanism, we consider the scenario that in an electronic marketplace a buyer 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).[3] 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:

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

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

The central server forwards the request to sellers in the marketplace. Sellers $\overline{S}$ that are interested in selling the product to the buyer can submit their bids containing their setting

[2] 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.

[3] 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.

for prices of the product, as well as values for non-price features. The buyer B will then determine the winning seller of the auction, which it can do business with.

In the sections that follow, we first describe how social network of buyers can be created by using our personalized approach. We then formalize how a seller should bid for the buyer's request, by considering the reputation of the buyer modeled based on the social network topology. Finally, we formalize how a buyer should determine the winner of the auction.

Social Network of Buyers

Our mechanism allows the central server to maintain for each buyer a fixed number of neighbors from which the buyer can trust and ask advice about sellers' trustworthiness. The central server models the trust value a buyer has of another buyer (an advisor) through a personalized approach. We first represent private reputation values, based on what is known about the advisors' ratings for sellers with which the buyer has already had some experience. Next, we describe how to construct a public model of trustworthiness of advisors based on common, centrally held knowledge of sellers and the ratings provided by advisors, including the trust ratings of sellers totally unknown to the buyer. We then outline how both private and public models can be combined, in order to obtain a value for the trustworthiness of each possible advisor.

The personalized approach[4] allows the central server to evaluate the private reputation the buyer B has of an advisor A by comparing their ratings for commonly rated sellers $\{S_1, S_2, ..., S_l\}$. For one of the commonly rated sellers S_i ($1 \leq i \leq l$ and $l \geq 1$), A has the rating vector $\overline{r_{A,S_i}}$ and B has the rating vector $\overline{r_{B,S_i}}$. A rating for S_i from B and A is binary ("1" or "0", for example), in which "1" means that the seller delivers the product and the valuation of the product is not less than that described in its bid, and "0" otherwise.[5] 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 $\overline{r_{A,S_i}}$ and $\overline{r_{B,S_i}}$ are ordered according to the time when they are provided. The ratings are then partitioned into different elemental time windows. The length of an elemental time window may be fixed (e.g. one day) or adapted by the frequency of the ratings to the seller S_i, similar to the way proposed in (Dellarocas 2000). It should also be considerably small so that there is no need to worry about the changes of sellers' behavior within each elemental time window. We define a pair of ratings (r_{A,S_i}, r_{B,S_i}), such that r_{A,S_i} is one of the ratings of $\overline{r_{A,S_i}}$, r_{B,S_i} is one of the ratings of $\overline{r_{B,S_i}}$, and r_{A,S_i} corresponds to r_{B,S_i}. The two ratings, r_{A,S_i} and r_{B,S_i}, are correspondent only if they are in the same elemental time window, the rating r_{B,S_i} is the most recent rating in its time window, and the rating r_{A,S_i} is the closest and prior to the rating r_{B,S_i}.[6] We then count the number of such pairs for S_i, N_{S_i}. The total number of rating pairs for all commonly rated sellers, N_{all} will be calculated by summing up the number of rating pairs for each commonly rated seller as follows:

$$N_{all} = \sum_{i=1}^{l} N_{S_i} \tag{2}$$

The private reputation of the advisor is estimated by examining rating pairs for all commonly rated sellers. We define a rating pair (r_{A,S_i}, r_{B,S_i}) as a positive pair if r_{A,S_i} is the same value as r_{B,S_i}. Otherwise, the pair is a negative pair. Suppose there are N_f number of positive pairs. The number of negative pairs will be $N_{all} - N_f$. The private reputation of the advisor A is estimated as the probability that A will provide reliable ratings to B. Because there is only incomplete information about the advisor, the best way of estimating the probability is to use the expected value of the probability. The expected value of a continuous random variable is dependent on a probability density function, which is used to model the probability that a variable will have a certain value. Because of its flexibility and the fact that it is the conjugate prior for distributions of binary events, the beta family of probability density functions is commonly used to represent probability distributions of binary events(see, e.g. the generalized trust models BRS (Jøsang & Ismail 2002) and TRAVOS (Teacy *et al.* 2005)). Therefore, the private reputation of A can be calculated as follows:

$$\alpha = N_f + 1, \beta = N_{all} - N_f + 1$$
$$R_{pri}(A) = E(Pr(A)) = \frac{\alpha}{\alpha + \beta}, \tag{3}$$

where $Pr(A)$ is the probability that A will provide fair ratings to B, and $E(Pr(A))$ is the expected value of the probability.

When there are not enough rating pairs, A's public reputation will also be considered. The public reputation of A is estimated based on its ratings and other ratings for the sellers rated by A. Each time A provides a rating $r_{A,S}$, the rating will be judged centrally as a fair or unfair rating. We define a rating for a seller as a fair rating if it is consistent with the majority of the ratings of the seller.[7] We consider only the ratings that are within the same time window as $r_{A,S}$, and we only consider the most recent rating from each advisor within any time window. In so doing, as sellers change their behavior and become more or less trustworthy to each advisor, the majority of ratings will be able to change.

[4]This approach was first introduced in (Zhang & Cohen 2006).

[5]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 would be used to represent probability distributions of ratings.

[6]We consider ratings provided by B after those by A in the same time window, in order to incorporate into B's rating anything learned from A during that time window, before taking an action. According to the solution proposed by Zacharia et al. (Zacharia, Moukas, & Maes 1999), by keeping only the most recent ratings, we can avoid the issue of advisors' "flooding" the system.

[7]Determining consistency with the majority of ratings can be achieved in a variety of ways, for instance averaging all the ratings and seeing if that is close to the advisor's rating.

Suppose that the advisor A totally provides N_{all}^A ratings. If there are N_f^A number of fair ratings, the number of unfair ratings provided by A will be $N_{all}^A - N_f^A$. In a similar way as estimating the private reputation, the public reputation of the advisor A is estimated as the probability that A will provide fair ratings. It can be calculated as follows:

$$\alpha' = N_f^A + 1, \beta' = N_{all}^A - N_f^A + 1$$

$$R_{pub}(A) = \frac{\alpha'}{\alpha' + \beta'}, \quad (4)$$

which also indicates that the more the percentage of fair ratings advisor A provides, the more reputable it will be.

To estimate the trustworthiness of advisor A, we combine the private reputation and public reputation values together. The private reputation and public reputation values are assigned different weights. The weights are determined by the reliability of the estimated private reputation value.

We first determine the minimum number of rating pairs needed for B to be confident about the private reputation value it has of A. The Chernoff Bound theorem (Mui, Mohtashemi, & Halberstadt 2002) provides a bound for the probability that the estimation error of private reputation exceeds a threshold, given the number of rating pairs. Accordingly, the minimum number of pairs can be determined by an acceptable level of error and a confidence measurement as follows:

$$N_{min} = -\frac{1}{2\varepsilon^2} ln \frac{1-\gamma}{2}, \quad (5)$$

where ε is the maximal level of error that can be accepted by B, and γ is the confidence measure. If the total weight of all rating pairs is larger than or equal to N_{min}, buyer B will be confident about the private reputation value estimated based on its ratings and the advisor A's ratings for all commonly rated sellers. Otherwise, there are not enough rating pairs, the buyer will not be confident about the private reputation value, and it will then also consider public reputation. The reliability of the private reputation value can be measured as follows:

$$w = \begin{cases} \frac{N_{all}}{N_{min}} & \text{if } N_{all} < N_{min}; \\ 1 & \text{otherwise.} \end{cases} \quad (6)$$

The trust value of A will be calculated by combining the weighted private reputation and public reputation values as follows:

$$Tr(A) = wR_{pri}(A) + (1-w)R_{pub}(A) \quad (7)$$

It is obvious that the buyer will consider less the public reputation value when the private reputation value is more reliable. Note that when $w = 1$, the buyer relies only on private reputation.

For a new buyer, the central server randomly assigns to it some other buyers with high public reputation as candidates for its neighbors. The new buyer 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.

Seller Bidding for Buyer's Request

A seller $S \in \overline{S}$ that is interested in selling p to B submits a bid to the central server. It 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:

$$U(p) = P(p) - C(p) \quad (8)$$

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:

$$V'(p) = \sum_{i=1}^{m} w_i D(f_i) - C(p) \quad (9)$$

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 $F(V')$, to show the possible values for $V'(p)$.

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:

$$P^*(p) = C(p) + \frac{\int_{V_L - C_H}^{V'(p)} F(x)dx}{F(V')} \quad (10)$$

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.

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 10 then should be changed to be:

$$P^*(p) = C(p) + \frac{\int_{V_L - C_H}^{V'(p)} F(x)dx}{F(V')} - V_D(R) \quad (11)$$

where $V_D(R)$ is the valuation of discount for the buyer B

with reputation $R(B)$.[8] Comparing Equations 10 and 11, the bidding price of the seller in Equation 11 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.

As discussed earlier, our mechanism allows the central server to maintain for each buyer a list of neighbors that it trusts the most. A seller can then 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:

$$R(B) = \begin{cases} \frac{N_B}{\theta} & \text{if } N_B < \theta; \\ 1 & \text{otherwise.} \end{cases} \tag{12}$$

The value of θ depends on the total number of buyers in the marketplace [9]. 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. This problem can be addressed within a centralized architecture. In this case, the seller is allowed to model the trustworthiness of a buyer by checking its ratings provided to the central server by the buyer. If the buyer has provided unfair ratings for the seller, the buyer will be considered untrustworthy by the seller. The seller can maintain a trustworthy buyer list and not enter into auctions of untrustworthy buyers. Trustworthy buyers always provide fair ratings for the seller. Based on the assumption that a trustworthy buyer's neighbors are also likely trustworthy, the seller would then use 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, using Equation 12.

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:

$$S_{win} = arg \max_{S \in \overline{S}} V(p) \tag{13}$$

[8] The amount of discount offered to a buyer depends on how much future profit the seller can gain from conducting the current business with the buyer. We will formalize this discount in our future work after we have better insight into how best to model the expected future profit.

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

The buyer chooses the winner of the auction among only sellers that are considered to be trustworthy. As an important component of our proposed marketplace model, the buyer models trustworthiness of a seller by also using a personalized approach. It models private reputation of the seller 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. The use of forgetting and discounting factors is included in this part of the model, in an effort to provide a somewhat richer modeling of agents. This is an extension of the personalized method for modeling advisors described in the "Social Network of Buyers" section.

Suppose that B has the rating vector $\overline{r_{B,S}}$, which contains all the ratings provided by B for the seller S. 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 $N_{pos,i}^B$ and the number of negative ratings $N_{neg,i}^B$ in each time window T_i. The private reputation of the seller S can be estimated through the beta family of probability density functions as follows:

$$R_{pri}(S) = \frac{\sum_{i=1}^{n} N_{pos,i}^B \lambda^{i-1} + 1}{\sum_{i=1}^{n} (N_{pos,i}^B + N_{neg,i}^B) \lambda^{i-1} + 2} \tag{14}$$

where λ ($0 \le \lambda \le 1$)is a forgetting rate. The forgetting rate is also introduced by Jøsang and Ismail (Jøsang & Ismail 2002) to deal with possible changes of the seller's behavior over time because old ratings will be given less weight than more recent ones. Note that when $\lambda = 1$ there is no forgetting, and when $\lambda = 0$ only the ratings that are within the current time window T_1 will be considered.

If the buyer B does not have enough personal experience with the seller S, it will also consider ratings provided by its neighbors. The buyer 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. These ratings will be discounted based on the trustworthiness of the advisor, so that the ratings from less trustworthy advisors will carry less weight than ratings from more trustworthy ones.

Jøsang (Jøsang 2001) provides a mapping from beliefs defined by the Dempster-Shafer theory to the beta function as follows:

$$\begin{cases} b = \frac{N_{pos,i}^{A_j}}{N_{pos,i}^{A_j} + N_{neg,i}^{A_j} + 2} \\ d = \frac{N_{neg,i}^{A_j}}{N_{pos,i}^{A_j} + N_{neg,i}^{A_j} + 2} \\ u = \frac{2}{N_{pos,i}^{A_j} + N_{neg,i}^{A_j} + 2} \end{cases} \tag{15}$$

where b, d and u represent belief, disbelief and uncertainty parameters, respectively. In our case, b represents the probability that the proposition that the seller is trustworthy is true, and d represents the probability of the proposition is false. Note that $b + d + u = 1$ and $b, d, u \in [0, 1]$. As also pointed out in (Jøsang & Ismail 2002) and (Yu & Singh 2003), beliefs and disbeliefs can be directly discounted by the trustworthiness of the advisor as follows:

$$\begin{cases} b' = Tr(A_j)b \\ d' = Tr(A_j)d \end{cases} \tag{16}$$

From Equations 15 and 16, we then can derive a discounting function for the amount of ratings provided by the advisor A_j as follows:

$$\begin{cases} D^{A_j}_{pos,i} = \frac{2Tr(A_j)N^{A_j}_{pos,i}}{(1-Tr(A_j))(N^{A_j}_{pos,i}+N^{A_j}_{neg,i})+2} \\ D^{A_j}_{neg,i} = \frac{2Tr(A_j)N^{A_j}_{neg,i}}{(1-Tr(A_j))(N^{A_j}_{pos,i}+N^{A_j}_{neg,i})+2} \end{cases} \tag{17}$$

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

In the same way as estimating the private reputation, the public reputation of the seller S can be calculated as follows:

$$R_{pub}(S) = \frac{[\sum_{j=1}^{k}\sum_{i=1}^{n} D^{A_j}_{pos,i}\lambda^{i-1}] + 1}{[\sum_{j=1}^{k}\sum_{i=1}^{n}(D^{A_j}_{pos,i} + D^{A_j}_{neg,i})\lambda^{i-1}] + 2} \tag{18}$$

The ratings provided by the advisors will be also discounted by the forgetting factor λ.

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

$$Tr(S) = w'R_{pri}(S) + (1 - w')R_{pub}(S) \tag{19}$$

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

$$w' = \begin{cases} \frac{N^B_{all}}{N_{min}} & \text{if } N^B_{all} < N_{min}; \\ 1 & \text{otherwise.} \end{cases} \tag{20}$$

where N^B_{all} is the total number of ratings provided by B for the seller. 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, which can be determined based on Equation 5.

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' ($0 < \gamma' < \delta' < 1$).

If there are no trustworthy sellers submitting bids, the winner of the auction will be selected among the sellers with trust values that are between δ' and γ'. 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.

Examples

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

Buyer's Neighbor List

We first provide an example to demonstrate how the central server models trust values a buyer B has of other buyers and chooses the most trustworthy ones as B's neighbors. In this example, we assume that each buyer can have at most one neighbor.

Consider the case where there are three other buyers (advisors) A_x, A_y and A_z. Each of them has rated only the five sellers (S_1, S_2, S_3, S_4, and S_5). Table 1 lists the ratings provided by each advisor A_j ($j \in \{x, y, z\}$) for the five sellers. 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 each advisor provides at most one rating within each time window. We also assume that those are the only ratings provided by them.

Table 2: Ratings Provided by the Buyer B

T	T_1	T_2	T_3	T_4	T_5
S_1	1	1	1	1	1
S_2	1	1	1	1	-
S_3	1	1	1	-	-
S_4	1	1	-	-	-
S_5	1	-	-	-	-

As can be seen from Table 2, the buyer B has also provided some ratings for the five sellers. The buyer B might have not provided any rating for some sellers within some time window. For example, it has provided only one rating for the seller S_5, which is in the time window T_1. We assume that the ratings provided by B are after those provided by A_x, A_y and A_z if they are within the same time window.

We compare the ratings provided by A_x, A_y and A_z in Table 1 and ratings provided by B in Table 2. The buyer B has the same number of rating pairs with each advisor ($N_{all} = 15$). However, B has different numbers of N_f positive rating pairs with A_x, A_y and A_z, which are listed in Table 3. Accordingly, as can be seen from Table 3, the private reputation values of A_x, A_y and A_z are different, in which the private reputation value of A_x is the highest and that of A_z is the lowest. The result indicates that the advisor A_x is most likely to provide fair ratings and have similar preferences with the buyer B, whereas A_z most likely will lie and have different preferences with B.

According to Table 1, the total number of ratings provided by each advisor is the same ($N^{A_j}_{all} = 25$). We also count the number of fair ratings each advisor provides. A rating here is considered as a fair rating when it is consistent with the majority of ratings for the seller within a same time window. Consider the case where all of the five sellers are reputable and the majority of ratings are fair. In this case, a rating of "1" provided by an advisor will be considered as a fair rating, whereas a rating of "0" will be considered as an unfair rating. From the advisors' ratings listed in Table 1, we can

Table 1: Ratings of Sellers Provided by Advisors

A_j	A_x					A_y					A_z				
T	T_1	T_2	T_3	T_4	T_5	T_1	T_2	T_3	T_4	T_5	T_1	T_2	T_3	T_4	T_5
S_1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0
S_2	1	1	1	1	1	0	1	0	1	1	0	0	0	0	0
S_3	1	1	1	1	1	1	0	1	0	0	0	0	0	0	0
S_4	1	1	1	1	1	1	0	0	1	0	0	0	0	0	0
S_5	1	1	1	1	1	1	1	0	0	1	0	0	0	0	0

Table 3: Private and Public Reputation Values of Advisors

A_j	A_x	A_y	A_z
N_f	15	8	0
α	16	9	1
β	1	8	16
$R_{pri}(A_j)$	0.94	0.53	0.06
$N_f^{A_j}$	25	12	0
α'	26	13	1
β'	1	14	26
$R_{pub}(A_j)$	0.96	0.48	0.04

see that ratings provided by the advisor A_x are all fair, the advisor A_z always lies, and some of the ratings provided by the advisor A_y are unfair. Table 3 lists the number of fair ratings provided by each advisor and the corresponding public reputation value of it. From Table 3, it is clear that the advisor A_x is most likely to provide fair ratings, and the advisor A_z most likely will lie.

Table 4: Trustworthiness of Advisors

ε	0.1	0.15	0.2
N_{min}	115	51	29
w	0.13	0.29	0.52
$Tr(A_x)$	0.957	0.954	0.950
$Tr(A_y)$	0.487	0.495	0.506
$Tr(A_z)$	0.043	0.046	0.05

To combine private reputation and public reputation, the weight w should be determined. The value of w depends on the values of ε and γ, and the number of rating pairs N_{all}, which is the same for every advisor in our example. Suppose we have a fixed value, 0.8 for γ, which means that the confidence value should be no less than 0.8 in order for the buyer to be confident with the private reputation values of advisors. In this case, the larger the value of ε the buyer sets, the more confident it is with the private reputation values of advisors, which also means that the more weight it will put on the private reputation values. Table 4 lists different acceptable levels of errors, their correspondent weights of private reputation values, and different results of trust values. It clearly indicates that A_x is the most trustworthy. As a result, the buyer B will choose A_x as its neighbor. In the examples that follow, we set $\varepsilon = 0.2$ and $\gamma = 0.8$. The trustworthiness of A_x is then 0.95.

Buyer Choosing Winning Seller

We then use an example to demonstrate how the buyer B models trustworthiness of sellers by considering ratings of sellers provided by its neighbors, and how it selects the winning seller to do business with. 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 5. 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_6, S_7, S_8 and S_9 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_6, S_7, S_8 and S_9 can be calculated according to Equation 14 as follows:

$$R_{pri}(S_6|S_7|S_8|S_9) = \frac{0+1}{(0+0)+2} = 0.5$$

Table 6: Ratings of Sellers Provided by A_x

T	T_1	T_2	T_3	T_4	T_5
S_6	0	0	0	1	1
S_7	-	-	-	-	-
S_8	1	1	1	1	1
S_9	1	1	1	1	0

Table 7: Discounted Amount of Ratings of Sellers

T	T_1	T_2	T_3	T_4	T_5
$D^{A_x}_{pos,i}(S_1)$	0	0	0	0.93	0.93
$D^{A_x}_{neg,i}(S_1)$	0.93	0.93	0.93	0	0
$D^{A_x}_{pos,i}(S_2)$	0	0	0	0	0
$D^{A_x}_{neg,i}(S_2)$	0	0	0	0	0
$D^{A_x}_{pos,i}(S_3)$	0.93	0.93	0.93	0.93	0.93
$D^{A_x}_{neg,i}(S_3)$	0	0	0	0	0
$D^{A_x}_{pos,i}(S_4)$	0.93	0.93	0.93	0.93	0
$D^{A_x}_{neg,i}(S_4)$	0	0	0	0	0.93

The buyer B then considers ratings of sellers provided by its neighbor A_x. The ratings of the sellers provided by

Table 5: 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

the advisor A_x are listed in Table 6. Note that the advisor A_x does not have ratings for the seller S_7 because A_x has not done business with S_7. The amount of positive or negative ratings provided by A_x within each time window will be discounted by using Equation 17. The discounted amount of positive and negative ratings of sellers is listed in Table 7. For example, the discount amount of positive ratings of seller S_6 in time window T_4 is calculated to be 0.93.

In this example, we set λ to be 0.9, which means that the buyer B does not have much forgetting. According to Equation 18, the public reputation of the sellers can be calculated as follows:

$$R_{pub}(S_6) = \frac{\sum_{i=4}^{5} 0.93 * 0.9^{i-1} + 1}{\sum_{i=1}^{5} 0.93 * 0.9^{i-1} + 2} = 0.39$$

$$R_{pub}(S_7) = 0.5, \quad R_{pub}(S_8) = 0.83, \quad R_{pub}(S_9) = 0.72$$

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 19 as follows:

$$Tr(S_6) = 0 * 0.5 + (1 - 0) * 0.39 = 0.39$$

$$Tr(S_7) = 0.5, \quad Tr(S_8) = 0.83, \quad Tr(S_9) = 0.72$$

We set the threshold δ' to be 0.7. In this case, only the sellers S_8 and S_9 will be considered as trustworthy sellers by the buyer B.

We suppose that the sellers S_8 and S_9 may have different costs of producing the product p with certain features. The bid submitted by the seller S_8 specifies that S_8 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_9 specifies that S_9 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:

$$V(p, S_8) = 0.4 * 5 + 0.6 * 10 - 4 = 4, \quad V(p, S_9) = 1$$

The value of the product in the bid of S_9 is lower than that of S_8. Seller S_8 will be selected as the winner. Buyer B pays S_8 the price of 4. Later on, seller S_8 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_{10} 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 5. The seller S_{10} needs to decide how to bid for each buyer's request. It models the reputation of each buyer.

Table 8: 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 8. 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 12 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.

Table 9: 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

According to the reputation of each buyer, seller S_{10} specifies its bid for each buyer's request. The non-price and price features in each bid and profit that each buyer can gain are listed in Table 9. 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.

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 a different number (from 2 to 20) of requests. In our experiments, we assume that each buyer will submit a rating for each of its requests. Therefore, buyers that have larger number of requests will provide larger number of ratings. We also 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.

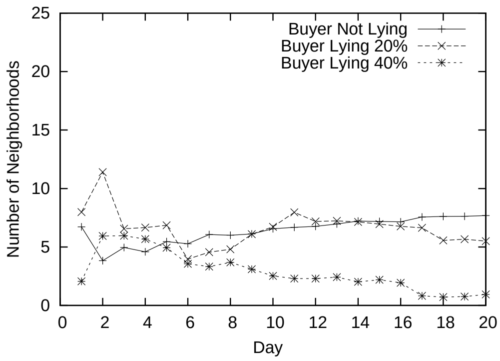

Figure 1: Reputation of Different Buyers

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 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 (that have provided more ratings) 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.

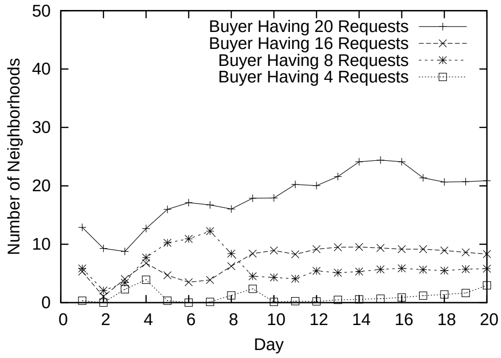

Figure 2: Reputation of Different Buyers

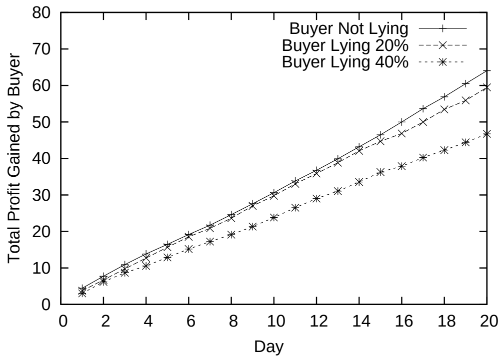

Figure 3: Profit Gained by Different Buyers

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 the valuation of the product received from its business partner. From Figure 3, we can see that buyers providing fewer unfair ratings will gain more total profit. Note that we do not measure total profit gained by buyers that have different numbers of requests. It is essential that the more requests the buyer has, the more profit it will be able to gain. In summary, it is better off for buyers to provide more fair ratings. Also note that the profit difference of different types of buyers is fairly small. It is because buyers have at most 20 requests in total.

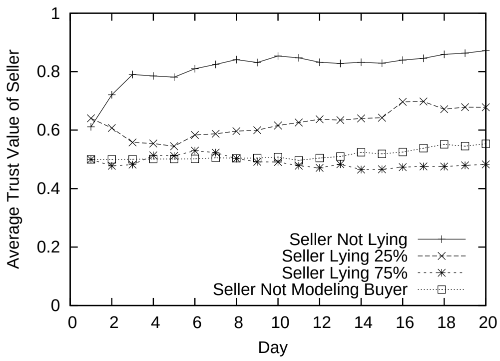

Figure 4: Average Trust Value of Different Sellers

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 (100 in our experiments). 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 and adjust their prices of products according to buyers' reputation 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_7 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.

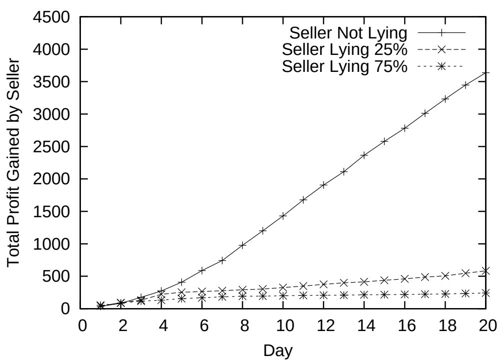

Figure 5: Total Profit Gained by Different Sellers

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. We can also see 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%. The reason is that we set the threshold δ' to be very high ($\delta' = 0.8$). The sellers lying 25% will not be considered as trustworthy sellers, therefore will have small chance to be selected as business partners by buyers. 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.

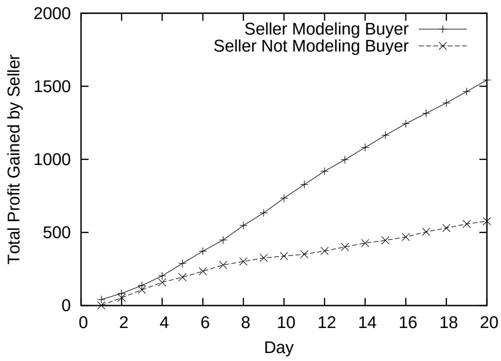

Figure 6: Total Profit Gained by Different Sellers

Related Work

There are other incentive mechanisms eliciting fair ratings. One type of such mechanisms is side payments (Dellarocas 2002; Jurca & Faltings 2003; Miller, Resnick, & Zeckhauser 2005). We survey three side payment mechanisms. They are different, for example, in terms of which party pays to honest buyers and/or in ways of evaluating the truthfulness of buyers' ratings. Another type of incentive mechanisms is credibility mechanisms (Papaioannou & Stamoulis 2005; Jurca & Faltings 2004) where only honest agents have their credibility in the marketplace enhanced. We point out some shortcomings of these methods and provide a contrast between our approach and those of other researchers.

Side Payment Mechanism

Dellarocas (Dellarocas 2002) proposes "Goodwill Hunting" (GWH) as a feedback mechanism for a trading environment based upon the argument that truthful feedback will benefit the community as a whole. This mechanism elicits truthful feedback from buyers by offering rebates of buyers' membership fee if the mean and variance between the buyers' and sellers' perception of quality of their transactions are consistent across the entire buyer community. Buyers may behave badly before they exit from the market. To solve this problem, part of the membership fee will be refunded only at the end of the period on the basis of buyers' behavior.

In the incentive compatible mechanism proposed by Jurca and Faltings (Jurca & Faltings 2003), a set of broker agents called R-agents, can sell and buy ratings of sellers to and from other ordinary agents. These ordinary agents first buy ratings from broker agents. After they finish doing business with sellers, they can sell ratings of the sellers back to the broker agents from which they bought ratings. To balance payoffs, ordinary agents are only allowed to sell ratings of a seller if they have previously bought reputation ratings of the seller. An agent will get paid only if a rating of a seller they provide is the same as the next rating of the same seller provided by another agent. A simple two agents case in an iterated Prisoner's Dilemma environment proves that the optimal strategy for an agent is to report truthfully because it will get paid with probability of at least 0.5.

Miller et al. (Miller, Resnick, & Zeckhauser 2005) introduces a mechanism which is very similar to that proposed by Jurca and Faltings (Jurca & Faltings 2003). In the mechanism, there is a center that maintains buyers' ratings. The center rewards or penalizes each buyer on the basis of its ratings and ensures that the mechanism at least breaks even in the long run. More specifically, a buyer providing truthful ratings will be rewarded and get paid not by broker agents but by the buyer after the next buyer. To balance transfers among agents, a proper scoring rule is used to determine the amount that each agent will be paid for providing truthful feedback. Scoring rules used by the center (i.e. the Logarithmic Scoring Rule) make truthful reporting a Nash equilibrium. Furthermore, proper scalings of scoring rules and collection of bonds or entry fees in advance ensure budget balance and incentives of the mechanism.

In summary, side payment mechanisms offer side payment to buyers that fairly rate results of business with sellers. However, they do not work well if the majority of buyers elect to provide unfair ratings because each of these dishonest buyers will receive a reward. This means that honest buyers that will not be giving similar ratings as many others, will not be rewarded and will be discouraged from being honest in the future. Moreover, this approach assumes that buyers act independently, and therefore has difficulty with the situation where buyers collude in giving unfair ratings. In contrast, in our mechanism, sellers can 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 limit their neighborhood of advisors to those that are determined to be trustworthy.

Credibility Mechanism

Instead of giving instant payment to agents that provide truthful ratings, credibility mechanisms measure agents' credibility or non-credibility according to their past ratings. It is believed that agents are more likely to conduct business with credible other ones.

One credibility mechanism is introduced by Papaioannou and Stamoulis (Papaioannou & Stamoulis 2005) for eliciting truthful ratings in peer-to-peer systems. Besides reputation information, each peer also stores a non-credibility value and a binary punishment state variable. After each transaction between two peers, they submit a rating indicating whether the transaction is successful or not. If both of them agree with the result of the transaction, their non-credibility values will be decreased. Otherwise, their non-credibility values will be increased and they will be punished. They will be forced not to conduct any transactions for a period determined by each of their non-credibility values.

A slightly different credibility mechanism called "CONFESS" is proposed by Jurca and Faltings (Jurca & Faltings 2004) for the online hotel booking industry. In this mechanism, a seller first reports its behavior. If it claims having cooperated, the buyer is then asked to submit a rating. If the buyer also reports that the seller has cooperated, it is sure that the seller has cooperated. Otherwise, both of them will be punished by decreasing their credibility as untruthful reporters because in this case at least one of them is cheating.

In credibility mechanisms, the credibility of two participants (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 these mechanisms, honest agents will be unfairly punished if they meet with a dishonest agent because they will not agree when they rate the results of their transactions with the agent. These honest agents will not gain credibility even if they provide good services. In addition, credibility mechanisms cannot deal with the situation where buyers and sellers collude to increase each other's credibility. Because our mechanism allows buyers to maintain a list of trustworthy other buyers as their neighbors, 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. In addition, all buyers have incentives to be honest, in order to enjoy the rewards offered by the honest sellers of the marketplace, if they maintain their position in many neighborhoods of the social network.

Conclusions and Future Work

In this paper, we propose a novel incentive mechanism to elicit fair ratings of sellers in electronic marketplaces. In our mechanism, a buyer maintains a neighbor list of other buyers that always provide fair ratings. We allow sellers to see how they have been rated by buyers and to model the reputation of buyers based on the social network created using our personalized approach. Reputable buyers are likely to be neighbors of many other buyers. Sellers then increase quality and/or decrease prices of 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. In summary, our mechanism is able to create a more effective electronic marketplace for buyers and sellers to do business with each other. In such an environment, honesty is promoted amongst buyers and sellers, and both honest parties participating in business are able to gain more profit. This also engenders trust of buying and selling agents from their human owners.

Our mechanism allows sellers to model reputation of a buyer based 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. We will also need to further study the properties of our social network, for example, the proper size of each neighbor list reflecting the population of buying and selling agents in the marketplace and how actively buying agents rate selling agents. Larger neighborhood size will increase the computation of maintaining and updating buying agents' neighbor lists, and may decrease the accuracy for predicting selling agents' trustworthiness from feedback provided by neighbors. Smaller neighborhood size may increase the accuracy, but will have higher chance the neighbors have insufficient experience (Herlocker, Konstan, & Riedl 2002).

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.

We will also develop more extensive experiments to validate our model. We are particularly interested in determining how robust our model is in coping with various types of collusion, including buyers colluding with sellers in providing unfairly high ratings and buyers colluding with other buyers in giving unfairly low ratings to sellers. It would also be useful to examine the case where some agents may vary their behavior widely, or where agents may enter and leave the marketplace.

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*.

Dellarocas, C. 2002. Goodwill hunting: An economically efficient online feedback mechanism for environments with variable product quality. *J. Padget and et al., editors, Agent-Mediated Electronic Commerce IV. Designing Mechanisms and Systems, volume LNCS 2531, Springer Verlag* 238–252.

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

Herlocker, J.; Konstan, J. A.; and Riedl, J. 2002. An empirical analysis of design choices in neighborhood-based collaborative filtering algorithms. *Information Retrieval* 5(4):287–310.

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

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

Jøsang, A. 2001. A logic for uncertain probabilities. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* 9(3):279–311.

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.

Mui, L.; Mohtashemi, M.; and Halberstadt, A. 2002. A computational model of trust and reputation. In *Proceedings of the 35th Hawaii International Conference on System Science (HICSS)*.

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)*.

Sabater, J., and Sierra, C. 2005. Review on computational trust and reputation models. *Artificial Intelligence Review* 24(1):33–60.

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

Teacy, W. T. L.; Patel, J.; Jennings, N. R.; and Luck, M. 2005. Coping with inaccurate reputation sources: Experimental analysis of a probabilistic trust model. In *Proceedings of 4th International Autonomous Agents and Multi Agent Systems (AAMAS 2005)*.

Yu, B., and Singh, M. 2003. Detecting deception in reputation management. In *Proceedings of International Autonomous Agents and Multi Agent Systems (AAMAS)*.

Zacharia, G.; Moukas, A.; and Maes, P. 1999. Collaborative reputation mechanisms in electronic marketplaces. In *Proceedings of the 32nd Hawaii International Conference on System Sciences (HICSS-32)*.

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*.