A Trust-Based Incentive Mechanism for E-Marketplaces

Jie Zhang, Robin Cohen, and Kate Larson

David R. Cheriton School of Computer Science
University of Waterloo
Waterloo, ON, Canada
{j44zhang,rcohen,klarson}@uwaterloo.ca

Abstract. In the context of electronic commerce, when modeling the trustworthiness of selling agent relies (in part) on propagating ratings provided by buying agents that have personal experience with the seller, the problem of unfair ratings arises. 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 trust-based 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.

1 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 on designing frameworks to model the trust and reputation of agents [1]. A modeling of the trustworthiness of a selling agent can be based on a buying agent's past personal experience with the seller (e.g. [2]). 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 (e.g. [3,4,5]). 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" [6]. 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".

Another problem is the fact that rating submission may be voluntary. Buyers may not have direct incentives to provide ratings because, for example, providing reputation ratings of sellers requires some effort [7,8]. Providing fair ratings for a trustworthy seller may also decrease the chance of doing business with the seller because of competition from other buyers.

R. Falcone et al. (Eds.): TRUST 2008, LNAI 5396, pp. 135–161, 2008.

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 [8,9], and credibility mechanisms [10,11]. 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) will be decreased if their ratings about the business result are different. As credibility values are made public, they will lose the chance to be selected as business partners. As a result, buyers are encouraged to provide fair ratings in order to keep up their credibility.

We instead propose to introduce a novel trust-based incentive mechanism into the e-marketplace. Our mechanism does not rely on side payment. Buyers in this mechanism are encouraged to be truthful in order to gain a greater number of profitable transactions. This idea is supported by work in the field of evolutionary game theory, such as the work of Gintis et al. [12]. 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 will benefit from its feedback through a greater number of profitable transactions.

More specifically, we first develop a model (a personalized approach) that allows buyers to effectively model the trustworthiness of other buyers(known as advisors) but with flexibility for buyers to weight how they value the contribution from the private reputation ratings and the public reputation ratings of these advisors.[1] A central element of our solution is to model the trustworthiness of advisors by tracking ratings according to the time when they are provided. In so doing, our approach is able to avoid the situation where advisors may untruthfully rate sellers a large number of times (known as "flooding") and is able to deal with changes in the behavior of agents.

We then use this personalized approach to create a social network of buyers. Each buyer in the society retains a neighborhood of the most trustworthy buyers, as advisors.[2] This social network is then leveraged by the sellers, as part of the promotion of honest reporting. In particular, sellers 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. Such buyers can then be considered reputable. This is also supported by Gintis et al. [12] 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 audience to witness their feedback (also known as increasing "broadcast efficiency"). The incentive for honesty is encouraged by having sellers in

[1] We use the term private reputation in the spirit of terms such as direct reputation used by researchers such as [13], to denote trustworthiness values obtained from personal information, rather than from indirect reports of trustworthiness.

[2] This approach is in line with that of researchers such as [14], used when constructing referral networks for information sharing.

our system increase quality and decrease prices of products to satisfy reputable buyers. This encourages 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.

2 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 the buyers and to determine their trustworthiness.

Consider the case where the buying and selling process is operated as a procurement (reverse) auction where the auctioneer is a buyer and bidders are sellers.[3] In this setting, 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 with the central server. Sellers that are interested in selling the product to the buyer will register to participate in the auction.

The buyer will first select the sellers it will consider for the auction, by modeling their trustworthiness. This is achieved by having each buyer maintain a neighborhood of trusted other buyers, which will be asked to provide ratings of the sellers under consideration. The buyer will then convey to the central server which sellers it is willing to consider, and the pool of possible sellers is thus reduced.

Sellers that are allowed to participate in the auction submit their bids that describe their settings for prices of the product and values of corresponding non-price features. The auction[4] is similar to Request For Quote (RFQ) introduced by Shachat and Swarthout [15], except that RFQ is an English auction and we propose using a first-price sealed auction for the purpose of saving communication costs of agents. Shachat and Swarthout in fact point out that an RFQ auction is equivalent to a first-price sealed bid auction.

[3] 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.

[4] Note that alternative auctions can also be deployed, such as English auction with Bidding Credits (EBC) [15]. However, the study of an alternative auction is outside the scope of this paper.

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 an amount equal to 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.

3 Incentive Mechanism

To formalize the proposed incentive mechanism, we consider the electronic marketplace scenario where 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 $\{f_1, f_2, ..., f_n\}$, as well as a set of weights $\{\omega_1, \omega_2, ..., \omega_n\}$ that correspond to each non-price feature.[5] 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).[6] We define the function $\tau()$ to denote such conversion. Inspired by [17], we also use a quasi-linear function to represent the buyer's valuation for the product as follows:

$$V_b = \sum_{i=1}^{n} \omega_i \tau(f_i) \tag{1}$$

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 register to participate in the auction.

In the sections that follow, we first describe how the social network of buyers can be created by using our personalized approach. We then formalize how a buyer should model sellers' trustworthiness and select the sellers it will consider for the auction, by asking for ratings of the sellers from its trusted neighbors. We also formalize how sellers that are allowed to join the buyer's auction should bid for the buyer's request by considering the reputation of the buyer modeled based on the social network topology. The buyer will finally choose the winner of the auction whose product (described in its bid) gives the buyer the largest profit, based on the buyer's evaluation criteria.

3.1 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.

[5] Note that buyers reveling their valuations to sellers is also used by [15] and [16].

[6] 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.

The central server records the trust value a buyer has of another buyer (an advisor) derived through a personalized approach. Buyers 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, buyers 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. Then both private and public models can be combined, in order to obtain a value for the trustworthiness of each possible advisor. Below we describe in detail how these calculations are done.[7]

In the personalized approach,[8] a buyer b may evaluate the private reputation it 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 $l \geq 1$), advisor a has the rating vector $\overline{r_{a,s_i}}$ and buyer 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), where "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.[9] 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 [6]. 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}.[10] 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}

7 The implementation of having a buyer truthfully compute another buyer's trustworthiness can be done by letting either a client-side application or the central server perform the computation.

8 This approach was first introduced in [18].

9 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.

10 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. [19], by keeping only the most recent ratings, we can avoid the issue of advisors' "flooding" the system.

is the same value as r_{b,s_i}. Otherwise, the pair is a negative pair. Suppose there are N_p positive pairs. The number of negative pairs will be $N_{all} - N_p$. 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 [3] and TRAVOS [4]). Therefore, the private reputation of a can be calculated as follows:

$$\alpha = N_f + 1, \quad \beta = N_{all} - N_p + 1$$

$$R_{pri}(a) = E[Pr(a)] = \frac{\alpha}{\alpha + \beta}, \tag{3}$$

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, advisor 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.[11] 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.

Suppose that the advisor a provides N'_{all} ratings in total. If there are N_f number of fair ratings, the number of unfair ratings provided by a will be $N'_{all} - N_f$. 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:

$$\alpha' = N_f + 1, \quad \beta' = N'_{all} - N_f + 1$$

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

which also indicates that the greater 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.

[11] 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.

We first determine the minimum number of rating pairs needed for buyer b to be confident about the private reputation value it has of advisor a. The Chernoff Bound Theorem [20] 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:

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

where $\varepsilon \in (0, 1)$ is the maximal level of error that will be accepted by b and $\gamma \in (0, 1)$ is the level of confidence buyer b would like to attain. If the total number 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:

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

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

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

It is obvious that the buyer will put less weight on 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.

3.2 Buyer Limiting Seller's Participation and Choosing a Winning Seller

To avoid doing business with possibly dishonest sellers, only sellers that are considered to be trustworthy by a buyer are allowed to join that buyer's auction. 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 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 also included in this part of the model.

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 $N^b_{pos,i}$ and the number of negative ratings $N^b_{neg,i}$ in each time window T_i. The private reputation of the seller s can be estimated through the beta family of probability density functions as follows:

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

where λ ($0 \leq \lambda \leq 1$) is a forgetting rate. The forgetting rate is also introduced by Jøsang and Ismail [3] 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 ($1 \leq j \leq k$) has provided $N^{a_j}_{pos,i}$ positive ratings and $N^{a_j}_{neg,i}$ 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 [21] provides a mapping from beliefs defined by the Dempster-Shafer theory to the beta function as follows:

$$\begin{cases} e = \frac{N^{a_j}_{pos,i}}{N^{a_j}_{pos,i}+N^{a_j}_{neg,i}+2} \\ d = \frac{N^{a_j}_{neg,i}}{N^{a_j}_{pos,i}+N^{a_j}_{neg,i}+2} \\ u = \frac{2}{N^{a_j}_{pos,i}+N^{a_j}_{neg,i}+2} \end{cases} \tag{9}$$

where e, d and u represent belief, disbelief and uncertainty parameters, respectively. In our case, e 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 $e+d+u = 1$ and $e, d, u \in [0, 1]$. As also pointed out in [3] and [5], beliefs and disbeliefs can be directly discounted by the trustworthiness of the advisor as follows:

$$\begin{cases} e' = Tr(a_j)e \\ d' = Tr(a_j)d \end{cases} \tag{10}$$

From Equations 9 and 10, we then can derive a discounting function for the amount of ratings provided by the advisor a_j as follows:

$$\begin{cases} D^{a_j}_{pos,i} = \frac{2Tr(a_j)\ N^{a_j}_{pos,i}}{(1-Tr(a_j))\ (N^{a_j}_{pos,i}+N^{a_j}_{neg,i})\ +\ 2} \\ D^{a_j}_{neg,i} = \frac{2Tr(a_j)\ N^{a_j}_{neg,i}}{(1-Tr(a_j))\ (N^{a_j}_{pos,i}+N^{a_j}_{neg,i})\ +\ 2} \end{cases} \tag{11}$$

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:

$$R_{pub}(s) = \frac{[\sum_{j=1}^{k}\sum_{i=1}^{n} D_{pos,i}^{a_j}\lambda^{i-1}] + 1}{[\sum_{j=1}^{k}\sum_{i=1}^{n}(D_{pos,i}^{a_j} + D_{neg,i}^{a_j})\lambda^{i-1}] + 2} \tag{12}$$

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:

$$Tr(s) = w' R_{pri}(s) + (1 - w')R_{pub}(s) \tag{13}$$

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

$$w' = \begin{cases} \frac{N_{all}^b}{N_{min}} & \text{if } N_{all}^b < N_{min}; \\ 1 & \text{otherwise.} \end{cases} \tag{14}$$

where N_{all}^b 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 θ ($0 < \theta < \delta < 1$).

If there are no trustworthy sellers register to the auction, the sellers with trust values that are between δ and θ will also be allowed to join the auction. Our idea of selective tendering is also supported by Kim's investigation results demonstrated in [22]. 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.

A set of m sellers allowed to participate in the auction will submit their bids. The way of formalizing their bids will be presented in the next section. 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:

$$s_{win} = \arg\max_{s_i, i=1}^{m} V_b \tag{15}$$

3.3 Seller's Bidding Behavior

A seller s_i ($1 \leq i \leq m$) that is allowed to join the buyer's auction sets the price and values for the non-price features of the product p, depending on how much instant 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 (labelled U, or utility) as follows:

$$U_{s_i} = P_{s_i} - C_{s_i} \tag{16}$$

where P_{s_i} is the price of the product set by the seller s_i and C_{s_i} is the cost for the seller to produce the product p with certain values for the non-price features in its bid.

The profit gained by the buyer if it chooses to do business with the seller s_i can be formalized as follows:

$$U_b = V_b - P_{s_i} \tag{17}$$

where V_b is determined using Equation 1. The buyer's profit is offered by the seller and may also be called the seller's "surplus offer", denoted as O_{s_i}. The seller s_i will try to gain profit from the transaction. It is reasonable to assume that $P_{s_i} \geq C_{s_i}$. Therefore, the best potential gain of the buyer from the transaction is when the price of the product is the same as the cost for the seller to produce the product, which can be formalized as follows:

$$S_{s_i} = V_b - C_{s_i} \tag{18}$$

S_{s_i} is so called "realized surplus", the best possible surplus for the buyer that the seller can offer. We also define the cumulative distribution function for S_{s_i} as $F()$ and the support of $F()$ is $[S_L, S_H]$. We assume $S_L \geq 0$ to ensure that the value of a seller's product always exceeds its cost.

The seller whose surplus offer, O_{s_i} is the highest will win the auction. The RFQ auction then becomes a first-price sealed auction. As argued by Shachat and Swarthout [15], a symmetric Bayes-Nash equilibrium surplus offer function can be derived as follows:

$$O^*_{s_i} = S_{s_i} - \frac{\int_{S_L}^{S_{s_i}} [F(x)]^{m-1} dx}{[F(S_{s_i})]^{m-1}} \tag{19}$$

where m is the number of bidders. Recall that O_{s_i} is the same as U_b. From Equations 17, 18 and 19, the equilibrium bidding function for the seller can then be derived as follows:

$$P^*_{s_i} = C_{s_i} + \frac{\int_{S_L}^{S_{s_i}} [F(x)]^{m-1} dx}{[F(S_{s_i})]^{m-1}} \tag{20}$$

The seller in our mechanism also reasons about the expected future gain from winning the current auction. It takes into account the reputation of the buyer b. In our mechanism, each buyer in the marketplace has a fixed number of neighbors that the buyer trusts the most and from which it can ask advice about sellers. This forms a social network of buyers. A buyer is reputable if it is the neighbor 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 to gain more profit in the future. We use R_b (reputation of b) to denote the number of other buyers considering b as their neighbor and $E_{s_i}(R_b)$ to denote the amount of the expected future gain. We then have the following inequality:

$$\frac{\partial[E_{s_i}(R_b)]}{\partial R_b} \geq 0 \tag{21}$$

Let us consider a scenario where sellers $\{s_1, s_2, ..., s_m\}$ have the same productivity. They have the same cost for producing the products that are valued equally by the buyer. In other words, we make the assumption that the distribution of S_{s_i}, $F()$ is a uniform distribution. Let us also assume that the seller's lowest realized surplus S_L for a transaction is 0. Equation 20 then can be simplified as follows:

$$\begin{aligned} P^*_{s_i} &= C_{s_i} + \frac{\int_{S_L}^{S_{s_i}} [F(x)]^{m-1} dx}{[F(S_{s_i})]^{m-1}} \\ &= C_{s_i} + \frac{\int_0^{S_{s_i}} (\frac{x}{S_H})^{m-1} dx}{(\frac{S_{s_i}}{S_H})^{m-1}} \\ &= C_{s_i} + \frac{S_{s_i}}{m} \end{aligned} \tag{22}$$

From Equations 16, 17 and 18, we can see that the seller's realized surplus is in fact equal to the sum of the buyer and the seller's profit. Since the seller has expected future gain from winning the current auction, the seller's realized surplus S_{s_i} can then be rewritten as follows:

$$\begin{aligned} S'_{s_i} &= U_b + U_{s_i} + \lambda' E_{s_i}(R_b) \\ &= V_b - C_{s_i} + \lambda' E_{s_i}(R_b) \\ &= S_{s_i} + \lambda' E_{s_i}(R_b) \end{aligned} \tag{23}$$

where $\lambda' \in [0, 1]$ is a discounting factor.[12] The lowest S'_{s_i} becomes $\lambda' E_{s_i}(R_b)$ instead of zero and the upper bound of S'_{s_i} becomes $S_H + \lambda' E_{s_i}(R_b)$. Accordingly, the symmetric Bayes-Nash equilibrium surplus offer function formalized in Equation 19 should be rewritten as follows:

$$O^*_{s_i} = S_{s_i} + \lambda' E_{s_i} - \frac{\int_{\lambda' E_{s_i}}^{S'_{s_i}} [F(x)]^{m-1} dx}{[F(S'_{s_i})]^{m-1}} \tag{24}$$

From Equations 17, 18 and 24, we then can derive the modified equilibrium bidding function for the seller as follows:

$$\begin{aligned} P^*_{s_i} &= C_{s_i} - \lambda' E_{s_i} + \frac{\int_{\lambda' E_{s_i}}^{S'_{s_i}} [F(x)]^{m-1} dx}{[F(S'_{s_i})]^{m-1}} \\ &= C_{s_i} - \lambda' E_{s_i} + \frac{\int_{\lambda E_{s_i}}^{S_{s_i}+\lambda' E_{s_i}} (\frac{x}{S_H})^{m-1} dx}{(\frac{S_{s_i}+\lambda' E_{s_i}}{S_H})^{m-1}} \\ &= C_{s_i} + \frac{S_{s_i}}{m} - \underbrace{\frac{1}{m}[\frac{(\lambda' E_{s_i})^m}{(S_{s_i} + \lambda' E_{s_i})^{m-1}} + (m-1)\lambda' E_{s_i}]}_{D_{s_i}} \end{aligned} \tag{25}$$

[12] The discounting factor is used to allow sellers to learn over time the likelihood of receiving their expected future gain.

We have already shown that sellers will gain better future profit when successful with more reputable buyers (Equation 21); this therefore suggests that the seller should offer better rewards to more reputable buyers as well. The bidding function outlined in Equation 25 provides for this, as the final term D_{s_i} in the equation becomes a positive term times the change in E_{s_i}. Note that since the value of $P^*_{s_i}$ in Equation 25 is smaller than that of Equation 22 the reward for buyers can either be lower price or higher cost (i.e. greater product quality). The seller sacrifices its current profit in both cases.

So far we have assumed that buyers do not attempt to collude with other buyers. However, there may be situations 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 with untrustworthy buyers. Trustworthy buyers 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.[13] From the list of all possible trustworthy buyers that the seller can find, the seller then can correctly model the reputation of a buyer. The robustness of our mechanism in coping with various types of collusion will be further discussed in Section 6 when comparing with other existing incentive mechanisms.

4 Examples

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

4.1 Buyer's Neighbor List

We first provide an example to demonstrate how buyer b models the trustworthiness of other buyers and chooses the most trustworthy ones as its neighbors. In this example, we assume that each buyer can have at most one neighbor.

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

[13] Recall that a neighbor of a buyer is trustworthy if it has produced like-minded ratings in the past.

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	-	-	-	-

Table 3. Private and Public Reputation Values of Advisors

a_j	a_x	a_y	a_z
N_p	15	8	0
α	16	9	1
β	1	8	16
$R_{pri}(a_j)$	0.94	0.53	0.06
N_f	25	12	0
α'	26	13	1
β'	1	14	26
$R_{pub}(a_j)$	0.96	0.48	0.04

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.

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_p 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, whereas a_z most likely will lie.

According to Table 1, the total number of ratings provided by each advisor is the same ($N'_{all} = 25$). 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 the 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 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.

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

4.2 Buyer Choosing a Winning Seller

We next 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 allows trustworthy ones to join its auctions. Suppose that the four sellers s_6, s_7, s_8 and s_9 are

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
$D^{a_x}_{pos,i}(s_6)$	0	0	0	0.93	0.93
$D^{a_x}_{neg,i}(s_6)$	0.93	0.93	0.93	0	0
$D^{a_x}_{pos,i}(s_7)$	0	0	0	0	0
$D^{a_x}_{neg,i}(s_7)$	0	0	0	0	0
$D^{a_x}_{pos,i}(s_8)$	0.93	0.93	0.93	0.93	0.93
$D^{a_x}_{neg,i}(s_8)$	0	0	0	0	0
$D^{a_x}_{pos,i}(s_9)$	0.93	0.93	0.93	0.93	0
$D^{a_x}_{neg,i}(s_9)$	0	0	0	0	0.93

all willing to sell the buyer the product p and register to participate in the auction. 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 8 as follows:

$$R_{pri}(s_6|s_7|s_8|s_9) = \frac{0+1}{(0+0)+2} = 0.5$$

The buyer b then considers ratings of sellers provided by its neighbor a_x. The ratings of the sellers provided by 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 11. 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 12, 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 13 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. Only these two sellers will be allowed to join the buyer's auction. They then submit their bids to 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.

4.3 Seller Bidding for Buyers' Requests

In this example, we illustrate how a seller s 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 the same evaluation criteria presented in Table 5, which specifies the two non-price features of p, the weight for each non-price feature and the information about the conversion from descriptive values to numeric values.

Seller s needs to decide how to bid for each buyer's request. It models the reputation of each buyer. Assume that each buyer is allowed to have only 3 neighbors in this

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

example. The neighbors of each buyer are listed in Table 8. We calculate each buyer's reputation represented by the number of its neighborhoods as follows:

$$R_{b_1} = 0, \quad R_{b_2} = 1, \quad R_{b_3} = 3$$

$$R_{b_4} = 4, \quad R_{b_5} = 5, \quad R_{b_6} = 5$$

According to the reputation of each buyer, seller s specifies its bid for each buyer's request. It produces different instantiations of the product p for different buyers. Table 9 lists the buyers' values for the products, calculated using Equation 1 based on Table 5. The seller s has different costs for producing these products, which are also listed in Table 9.

Table 9. Products Produced for Different Buyers

Buyers	Non-price Features		Value	Cost
	Delivery Time	Warranty		
b_1, b_2	7 days	1 year	3	1
b_3, b_4	3 days	2 years	5	3
b_5, b_6	1 day	3 years	10	8

Table 10 lists the seller's amount of expected future gain $E_s(R_b)$ from selling the products to the buyers with different reputation values. We assume the discounting factor λ' to be 1 and the number of bidders in each auction is 5. We also calculate the realized surplus S_s using Equation 23, and the reward D_s offered to different buyers and the seller's equilibrium bidding prices P_s^* according to Equation 25, as presented in Table 10.

Table 10. Seller's Prices for Different Buyers

Buyer	b_1	b_2	b_3	b_4	b_5	b_6
$E_s(R_b)$	0	0.1	0.3	0.4	0.5	0.5
S_s	2	2.1	2.3	2.4	2.5	2.5
D_s	0	0.08	0.24	0.32	0.4	0.4
P_s^*	1.4	1.34	3.22	3.16	8.09	8.09

We can see from Table 10 that seller s offers the best rewards to the more reputable buyers b_5 and b_6. Buyer b_1 with reputation value 0 does not gain any reward. According to Tables 9 and 10, we can calculate the profit gained by the buyers using Equation 17, as follows:

$$U_{b_1} = 1.6, \quad U_{b_2} = 1.66, \quad U_{b_3} = 1.78$$

$$U_{b_4} = 1.84, \quad U_{b_5} = 1.91, \quad U_{b_6} = 1.91$$

We can see that the more reputable buyers b_5 and b_6 are able to gain the largest profit and the less reputable buyers b_1 and b_2 can only gain the smallest profit. Therefore, it is better off for buyers to be honest and build higher reputations, in order to gain more profit.

5 Experimental Results

This section presents experimental results to confirm the value of our proposed incentive mechanism, showing that: honesty is more profitable, for both buyers and sellers; sellers are more profitable when modeling the reputation of buyers according to their neighborhoods; buyers are more profitable when they participate, by providing ratings to others; buyers derive better profit when they use the ratings of sellers provided by neighbors and measure the trustworthiness of other buyers, in order to form these neighborhoods.

We simulate a marketplace operating with our mechanism for a period of 30 days. The marketplace involves 90 buyers. These buyers are grouped into three groups. They have different numbers of requests. Every 10 of the buyers in each group has a different number (10, 20 and 30) of requests. In our experiments, we assume that there is only one product in each request and each buyer has a maximum of one request each day. For the purpose of simplicity, we also assume that the products requested by buyers have the same non-price features. After they finish business with sellers, buyers rate sellers. Some buyers will provide unfair ratings. Each group of buyers provides different percentages (0%, 20% and 40%) of unfair ratings. We allow 2 buyers from each group to leave the marketplace at the end of each day. Accordingly, we also allow 6 buyers to join the marketplace at the end of each day. These buyers will also provide different percentage (0%, 20% and 40%) of unfair ratings, to keep the number of buyers in each group the same. Initially, we randomly assign 5 buyers to each buyer as its neighbors.

There are also 9 sellers in total in the marketplace. Each 3 sellers acts dishonestly in different percentages (0%, 25% and 75%) of their business with buyers. We assume that all sellers have the same cost for producing the products because all products have the same non-price features.

5.1 Promoting Honesty

Here, we provide some general results to show that our mechanism promotes buyer and seller honesty. We first measure the reputation of buyers that provide different percentages of unfair ratings. In our experiments, a buyer's reputation is represented by the number of other buyers considering this buyer as their neighbor. The results[14] are shown in Figure 1. From this figure, we can see that the buyers providing the smaller percentages of unfair ratings will have the larger reputation values. Due to the randomness of the initial setting for our experiments, buyers' reputation values change stochastically at the beginning. After approximately 6 days when our marketplace converges, the changes of buyers' reputation will clearly follow a trend.

After each day, we measure total profit gained by buyers that provide different percentages of unfair ratings. The profit gained by a buyer from buying a product is formalized in Equation 17. From Figure 2, we can see that buyers providing fewer unfair ratings will gain more total profit. Note that even the profit difference of different types of buyers is fairly small because buyers have at most 30 requests in total in our experiment, we still can conclude that it is better off for buyers to provide truthful ratings of sellers.

[14] All experimental results in this section are averaged over 500 rounds of the simulation.

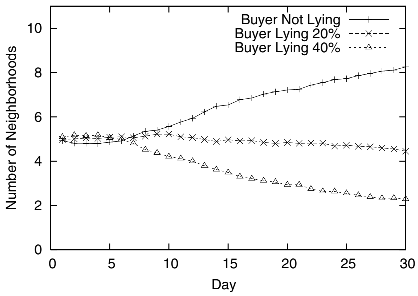

Fig. 1. Reputation of Different Buyers

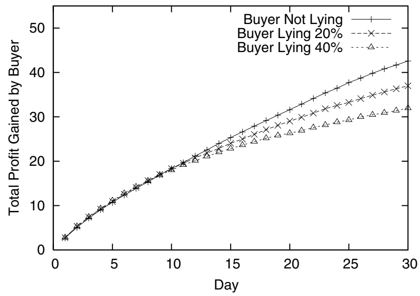

Fig. 2. Profit Gained by Different Buyers

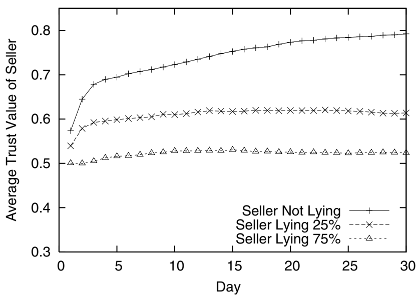

Fig. 3. Average Trust Value of Different Sellers

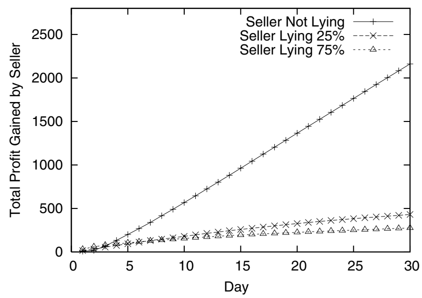

Fig. 4. Total Profit Gained by Different Sellers

We compare the average trust values of different sellers. The average trust value of a seller is calculated as the sum of the trust value each buyer has of the seller divided by the total number of buyers in the marketplace (90 in our experiments). As shown in Figure 3, results indicate that sellers being dishonest more often will have smaller average trust values. From this figure, we can see that the average trust values of the sellers being dishonest in 75% of their business are nearly 0.5.[15] This is because they do not have much chance to do business with buyers and will not have many ratings. A seller without any ratings will have a default trust value of 0.5.

We also compare total profit gained by different sellers. Results are shown in Figure 4. From this figure, 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% of the time is much larger than that between the sellers lying 25% and the sellers lying 75% of the time. The reason is that we set the threshold δ for sellers to be considered trustworthy to be very high. The sellers lying 25% of the time will not be considered as trustworthy sellers and therefore will have few occasions to be selected as business partners by buyers.

[15] Note that 25% of the time these sellers are honest and do gain some trust.

5.2 Seller Strategy

The purpose of this experiment is to examine the average trustworthiness of and the total profit gained by sellers using different strategies. We have two groups of sellers. One group of sellers will model reputation of buyers and offer better rewards to reputable buyers. Another group of sellers will not model reputation of buyers and ask for the same price from different buyers. Sellers in each group will lie in different percentages (0%, 25% and 75%) of their business with buyers. Buyers in this experiment adopt the strategy outlined in Section 3.2.

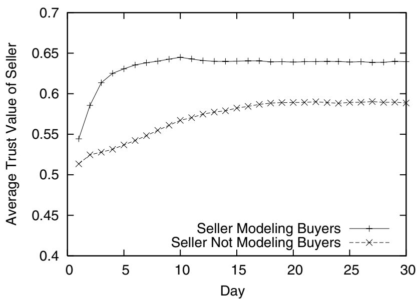

Fig. 5. Average Trust Value of Different Sellers

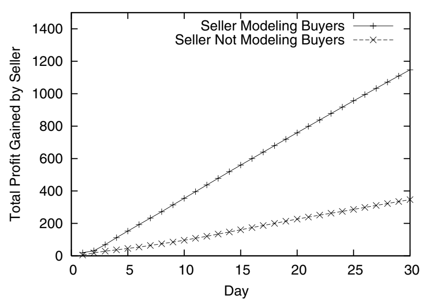

Fig. 6. Total Profit Gained by Different Sellers

We measure the average trust values of sellers from each group. Results shown in Figure 5 indicate that sellers modeling reputation of buyers will have higher average trust values. We also measure the total profit gained by different sellers. 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.

5.3 Buyer Strategy

Buyers in the marketplace may also have different strategies. They may not always provide ratings for sellers. They may use different methods to model sellers, or may not model others at all. In this section, we carry out experiments to compare reputation values and total profit of buyers using different strategies, while the sellers in the marketplace are using the strategy described in Section 3.3. Results show that our mechanism provides incentives for buyers to provide ratings of sellers and the modeling methods we propose will provide buyers with more profit.

Incentives for Providing Ratings. We examine the expectation of our mechanism that provides incentives for buyers to provide ratings. We compare reputation values and total profit of buyers providing different number of ratings. In this experiment, all buyers are honest. They have the same number of requests. However, they rate different percentages (1/3, 2/3 and 3/3) of their business with sellers.

We first measure the reputation of the buyers. Results are shown in Figure 7. Buyers that have provided more ratings will have larger reputation values. We also measure

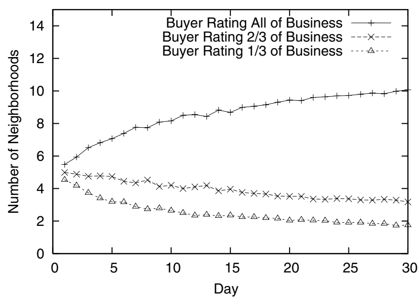

Fig. 7. Reputation of Different Buyers

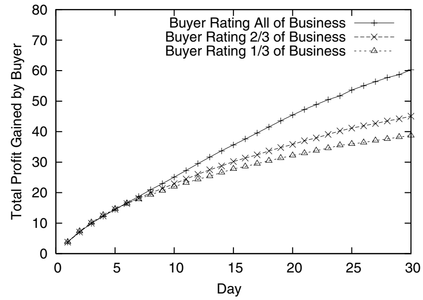

Fig. 8. Profit Gained by Different Buyers

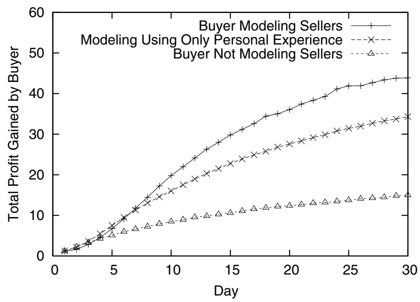

Fig. 9. Profit Gained by Different Buyers

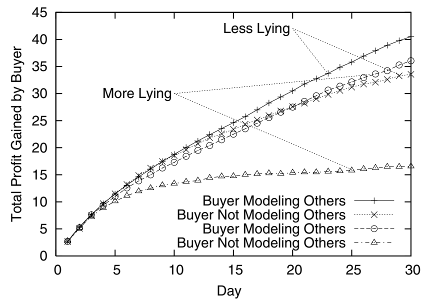

Fig. 10. Profit Gained by Different Buyers

total profit of these buyers. Results shown in Figure 8 indicate that buyers that have provided more ratings will be able to gain more total profit. Therefore, it is better off for buyers to provide ratings of sellers.

Buyer Modeling Sellers. In this experiment, one third of the buyers models the trustworthiness of sellers based on their personal experience with the sellers (private reputation of sellers) and advice about the sellers provided by their neighbors (public reputation of sellers). Another third of the buyers uses only personal experience to model the trustworthiness of sellers. These buyers select the most trustworthy sellers to join their auctions. The rest of the buyers do not model sellers. They allow every seller to submit a bid.

We compare the total profit gained by these three types of buyers. Results are shown in Figure 9. From this figure, we can see that buyers modeling the trustworthiness of sellers and selecting the most trustworthy sellers to participate in their auctions will be able to gain more total profit. It is also clear that buyers modeling sellers by taking into account as well the advice provided by other buyers will be able to gain more profit. In summary, it is better off for buyers to selectively choose sellers to participate in their auctions and to take into account the advice provided by other buyers when buyers lack personal experience with sellers.

Buyer Modeling Other Buyers. We have two different settings for this experiment. In the first setting, the first group of buyers does not provide any unfair ratings, but the second and third groups provide 20% and 40% of unfair ratings respectively. In the second setting, the first group of buyers still does not lie. The second and third groups lie more. They provide 50% and 100% of unfair ratings respectively. In both of the settings, one half of the buyers in the first group model other buyers and select the most trustworthy ones as their neighbors from which they can ask advice about sellers. Another half of the buyers do not model the trustworthiness of other buyers. They randomly select some other buyers as their neighbors.

We compare the total profit gained by these two types of buyers in the two settings. Results are shown in Figure 10. From this figure, we can see that buyers modeling the trustworthiness of other buyers and selecting the most trustworthy ones as their neighbors will be able to gain more total profit. It is also clear that the buyers that do not model the trustworthiness of other buyers will gain much less profit when the other buyers provide a lot of unfair ratings. Therefore, it is better off for buyers to model the trustworthiness of other buyers and select the most trustworthy ones as their neighbors from which they ask advice about sellers.

6 Related Work

There are other incentive mechanisms eliciting fair ratings. One type of such mechanisms is side payments [8,9,23]. 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 [10,11] 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.

6.1 Side Payment Mechanism

Dellarocas [23] 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 [9], 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. [8] introduces a mechanism which is very similar to that proposed by Jurca and Faltings [9]. 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. Jurca and Faltings [24] investigate mechanisms that can cope with collusion among buyers. However, they do not consider the case where a seller may collude with a group of buyers in promoting the seller itself or bad-mouthing another seller. 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.

6.2 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 [10] 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 [11] 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.

7 Conclusions and Future Work

In this paper, we propose a novel incentive mechanism to elicit fair ratings of sellers in electronic marketplaces that hinges on the social network of buyers. 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 incentive mechanism is built based on trust modeling and can further help trust modeling by diminishing the problem of unfair ratings.

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, using a method similar to the EigenTrust [25] and PageRank [26] computation. 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 the buying agents are rating 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 there will be a higher chance that the neighbors have insufficient experience [27]. Another avenue for future work is to examine more specific calculations for public reputation. Approaches such as [28] that combine trust beliefs from different buyers may be of some value.

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.

For future work, we will also look carefully into how sellers should model their expected future gain from winning the current auctions. The insights from research in viral marketing [29] about modeling buyers' influence in a market and sellers' profit increase from marketing to these buyers can be applied here. Another topic for future work is to determine the number of sellers allowed to join each buyer's auction, which ensures that dishonest sellers' instant profit does not exceed honest sellers' long-term profit. Kim [22] provides some insights into how to derive an optimal number of bidders.

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. In our future experiments, we could also examine the situation where selling agents may strategically vary their behavior to exploit the marketplace, which has been well studied by Sen and Banerjee [30].

References

1. Sabater, J., Sierra, C.: Review on computational trust and reputation models. Artificial Intelligence Review 24(1), 33–60 (2005)
2. Tran, T., Cohen, R.: Improving user satisfaction in agent-based electronic marketplaces by reputation modeling and adjustable product quality. In: Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2004) (2004)
3. Jøsang, A., Ismail, R.: The beta reputation system. In: Proceedings of the 15th Bled Electronic Commerce Conference (2002)

4. Teacy, W.T.L., Patel, J., Jennings, N.R., Luck, M.: Coping with inaccurate reputation sources: Experimental analysis of a probabilistic trust model. In: Proceedings of 4th International Autonomous Agents and Multiagent Systems (AAMAS) (2005)
5. Yu, B., Singh, M.: Detecting deception in reputation management. In: Proceedings of International Autonomous Agents and Multi Agent Systems (AAMAS) (2003)
6. Dellarocas, C.: Immunizing online reputation reporting systems against unfair ratings and discriminatory behavior. In: Proceedings of the 2nd ACM Conference on Electronic Commerce (EC), Minneapolis, MN (2000)
7. Jøsang, A., Ismail, R., Boyd, C.: A survey of trust and reputation systems for online service provision. Decision Support Systems (to appear, 2005)
8. Miller, N., Resnick, P., Zeckhauser, R.: Eliciting informative feedback: The peer-prediction method. Management Science 51(9), 1359–1373 (2005)
9. Jurca, R., Faltings, B.: An incentive compatible reputation mechanism. In: Proceedings of the IEEE Conference on E-Commerce (CEC), Newport Beach, CA, USA (2003)
10. Papaioannou, T.G., Stamoulis, G.D.: An incentives mechanism promoting truthful feedback in peer-to-peer systems. In: Proceedings of IEEE/ACM CCGRID 2005 (Workshop on Global P2P Computing (2005)
11. Jurca, R., Faltings, B.: 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 (2004)
12. Gintis, H., Smith, E.A., Bowles, S.: Costly signaling and cooperation. Journal of Theoretical Biology 213, 103–119 (2001)
13. Mui, L., Halberstadt, A., Mohtashemi, M.: Notions of reputation in multi-agent systems: A review. In: Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems (2002)
14. Yu, B., Venkatraman, M., Singh, M.P.: An adaptive social network for information access: Theoretical and experimental results. Applied Artificial Intelligence 17(1), 21–38 (2003)
15. Shachat, J., Swarthout, J.T.: Procurement auctions for differentiated goods. Working Paper, IBM Research Labs (2003)
16. Che, Y.K.: Design competition through multidimensional auctions. The RAND Journal of Economics 24(4), 668–680 (1993)
17. Boutilier, C., Sandholm, T., Shields, R.: Eliciting bid taker non-price preference in (combinatorial) auctions. In: Proceedings of the Nineteenth National Conference on Artificial Intelligence (AAAI 2004) (2004)
18. Zhang, J., Cohen, R.: A personalized approach to address unfair ratings in multiagent reputation systems. In: Proceedings of the 5th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS) Workshop on Trust in Agent Societies (2006)
19. Zacharia, G., Moukas, A., Maes, P.: Collaborative reputation mechanisms in electronic marketplaces. In: Proceedings of the 32nd Hawaii International Conference on System Sciences (HICSS-32), Maui, Hawaii (1999)
20. Mui, L., Mohtashemi, M., Halberstadt, A.: A computational model of trust and reputation. In: Proceedings of the 35th Hawaii International Conference on System Science (HICSS) (2002)
21. Jøsang, A.: A logic for uncertain probabilities. International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems 9(3), 279–311 (2001)
22. Kim, I.G.: A model of selective tendering: Does bidding competition deter opportunism by contractors? The Quarterly Review of Economics and Finance 38(4), 907–925 (1998)
23. Dellarocas, C.: Goodwill hunting: An economically efficient online feedback mechanism for environments with variable product quality. In: Padget, J., Shehory, O., Parkes, D.C., Sadeh, N.M., Walsh, W.E. (eds.) AMEC 2002. LNCS, vol. 2531, pp. 238–252. Springer, Heidelberg (2002)

24. Jurca, R., Faltings, B.: Collusion-resistant, incentive-compatible feedback payments. In: Proceedings of the 8th ACM Conference on Electronic Commerce (EC) (2007)
25. Kamvar, S.D., Schlosser, M.T., Garcia-Molina, H.: The eigentrust algorithm for reputation management in p2p networks. In: Proceedings of the 12th international conference on World Wide Web (WWW), pp. 640–651 (2003)
26. Page, L., Brin, S., Motwani, R., Winograd, T.: The PageRank Citation Ranking: Bringing order to the Web. Technical Report, Computer Science Department, Stanford University (1998)
27. Herlocker, J., Konstan, J.A., Riedl, J.: An empirical analysis of design choices in neighborhood-based collaborative filtering algorithms. Information Retrieval 5(4), 287–310 (2002)
28. Wang, Y., Singh, M.: Formal trust model for multiagent systems. In: Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI) (2007)
29. Richardson, M., Domingos, P.: Mining knowledge-sharing sites for viral marketing. In: Proceedings of the 8th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD) (2002)
30. Sen, S., Banerjee, D.: Monopolizing markets by exploiting trust. In: Proceedings of the 5th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS) (2006)