

Leveraging a Social Network of Trust for Promoting Honesty in E-Marketplaces

Jie Zhang*, Robin Cohen, and Kate Larson

*School of Computer Engineering, Nanyang Technological University, Singapore
School of Computer Science, University of Waterloo, Canada
{zhangj}@ntu.edu.sg

Abstract. In this paper, we examine a trust-based framework for promoting honesty in e-marketplaces that relies on buyers forming social networks to share reputation ratings of sellers and sellers rewarding the buyers that are most respected within their social networks. We explore how sellers reason about expected future profit when offering particular rewards for buyers. We theoretically prove that in a marketplace operating with our mechanism: i) buyers will be better off honestly reporting seller ratings and ii) sellers are better off being honest, to earn better profit. Experiments confirm the robustness of the approach, in dynamically changing environments. With rational agents preferring to be honest, the buyer and seller strategies as specified constitute an effective approach for the design of e-marketplaces.

1 Introduction

Artificial intelligence researchers have proposed the use of intelligent agents to act on behalf of buyers and of sellers, in electronic marketplaces. These agents are capable of learning over time the behavior of their business partners, to enable each party to make effective decisions about which parties they wish to do business with, in the future. One approach that has received much attention is to have buying agents model the trustworthiness of selling agents, making use of ratings of sellers provided by other buyers in the marketplace; this is of particular benefit when buyers do not have much personal experience with the sellers. However, the important problem where buyers may provide unfair ratings to bad-mouth some sellers or promote some other sellers has to be addressed [1].

In our previous work [2], trust modeling was promoted as an important avenue for creating incentives for buyer and seller honesty in the marketplace. In this approach, buying agents make use of a neighborhood of other buying agents (known as advisors) to provide seller ratings and sellers are inclined to act honestly, because they are being modeled. But sellers also offer better rewards to buyers that belong to many neighborhoods in the marketplace, a feature motivated by the work of Gintis et al. [3] which argues that altruism in one context signals “quality” that is rewarded by increased opportunities in other contexts. Experimental evidence was presented to demonstrate that a framework like this

promotes honesty in reporting seller ratings, from buyers, and honesty in delivering goods as promised, by sellers.

In this paper, we begin with the model of [2], where buyers use a social network to model sellers and sellers model buyers to offer varying rewards. We then explore how sellers can reason about expected future profit, starting from refined formulae for reasoning about immediate profit. As a result, we are able to provide a specification for seller bidding behavior and for offering rewards to buyers based on their reputation. We also emphasize the importance for buyers to adopt a strategy to limit the number of sellers that are considered for each good to be purchased. Most importantly, we theoretically prove that both rational buyers and rational sellers are incentivized to behave honestly in our mechanism, in so doing providing definitive validation of the effectiveness of our proposal.

We then present a series of experimental results to provide additional detail on marketplace trends that demonstrate the value of our newly designed incentive mechanism, conducted in a simulated environment where buyers and sellers may be deceptive and they may be arriving and departing. This provides a stronger endorsement of the mechanism as one that is robust to important conditions in the marketplace. In addition, we validate the benefit of our specific proposal for the seller bidding strategy and for the buyer strategy of limiting the sellers being considered, clearly showing the gains in profit enjoyed by both sellers and buyers when our mechanism is introduced and our proposed strategies are followed.

2 System Overview

The electronic marketplace environment we are modeling is populated with self-interested buying and selling agents. Our incentive mechanism is generally applicable to any marketplace where sellers may alter quality and price of their products to satisfy buyers. For the remainder of this paper, we discuss the scenario where the buyers and sellers are brought together by a procurement (reverse) auction, where the auctioneer is a buyer and bidders are sellers. There is a central server that runs the auction.

In our system, a buyer that wants to purchase a product sends a request to the central server. This request indicates not only the product but also the buyer's evaluation criteria for the product (discussed in more detail in the following section). Sellers interested in selling the product to the buyer will register to participate in the auction.

The buyer will first limit the sellers it will consider for the auction, by modeling their trustworthiness. This is achieved by having the central server retain for each buyer 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 allowed to participate in the auction will submit their bids and the buyer will select the winner of the auction as the seller whose product (described in its bid) gives the largest profit, based on the buyer’s evaluation criteria.

In order to formulate their bids, sellers model the reputation of buyers and make more attractive offers to more reputable buyers. A buyer’s reputation is based on the number of other buyers considering this buyer as their neighbor (as well as the trust these other buyers place on this buyer, and the reputation of these other buyers). The reputation of each buyer is maintained by the central server and released to the sellers.

Once a buyer has selected the winning seller, it pays that seller the amount indicated in the bid. The winning seller is supposed to deliver the product to the buyer. However, it may decide to alter the quality of the product or to not deliver the product at all. The buyer will report the result of conducting business with the seller to the central server, registering a rating for the seller. It is precisely these ratings of the seller that can then be shared with those buyers that consider this buyer as their neighbor.

In summary: the central server runs the auction and maintains information that is shared with sellers and buyers; buyers announce their intention to purchase products, consult with neighbors, choose a winning seller and report a final rating for the seller; sellers bid to win the sale to the buyer, consider buyer reputation in formulating their bids and then decide what product to deliver to the buyer (if at all).

3 Strategic Behavior Analysis

In this section, we propose and analyze the strategies that buyers and sellers in our mechanism should use. We also theoretically prove that these strategies will promote buyer and seller honesty.

3.1 Seller Strategy to Promote Buyer Honesty

We first present a seller’s optimal strategy when sellers only take into account their instant profit. We then derive an equilibrium bidding strategy for sellers when they also take into account their expected future gain, in a simplified scenario where all sellers have the same productivity. We then remove the simplifying assumption and show that with this bidding structure, sellers are better off providing rewards to more reputable buyers and that buyers are better off participating in the social network and providing honest ratings of sellers.

Seller Strategy We discuss our mechanism in the context of the Request For Quote (RFQ) system [4]. Building on the incentive mechanism outlined in [2], we consider a scenario where a buyer b wants to buy a product p . The buyer specifies its evaluation criteria for a set of non-price features $\{f_1, f_2, \dots, f_n\}$, as well as a set of weights $\{w_1, w_2, \dots, w_n\}$ 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, a 3-year warranty is converted to the numeric value of 10 on a scale of 1 to 10).¹ We define the function $\tau()$ to denote such a conversion. Sellers $\{s_1, s_2, \dots, s_m\}$ ($m \geq 1$) allowed to join the auction are able to know the buyer's values of their products, formalized as follows:

$$V_b = \sum_{j=1}^n w_j \tau(f_j) \quad (1)$$

Sellers will then need to formulate their bids, reflecting what they promise to deliver (e.g. a delivery time of 1 week). An honest seller is one that honors its promise.

As in the mechanism of [2], a seller s_i ($1 \leq i \leq m$) 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. The seller's instant profit is defined as follows:

$$U_{s_i} = P_{s_i} - C_{s_i} \quad (2)$$

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 seller s_i will try to gain profit from the transaction. It is reasonable to assume that $P_{s_i} \geq C_{s_i}$.

The elements above are introduced in the mechanism of [2]. We now extend the treatment, to express more precisely the profit to be gained by the buyer and the seller, to then discuss the kind of gains that sellers can reason about and the kind of bids they should offer to buyers (and to demonstrate how honesty is thus promoted in the marketplace).

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

$$U_b = V_b - P_{s_i} \quad (3)$$

The buyer's profit is also called the seller's "surplus offer", denoted as O_{s_i} . The seller's "realized surplus" is typically calculated as the sum of the buyer's and the seller's profit:

$$S_{s_i} = V_b - C_{s_i} \quad (4)$$

Note that the seller's realized surplus is higher when its cost C_{s_i} is lower. We also define the cumulative distribution function for S_{s_i} (over all sellers) 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.

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

The seller whose surplus offer is the highest will win the auction. The RFQ auction then becomes a first-price sealed auction where a bidder's bids are not seen by others and the bidder with the highest bid (surplus offer) wins the auction. As argued in [4], a symmetric Bayes-Nash equilibrium surplus offer function can be derived as follows:

$$O_{s_i}^* = S_{s_i} - \frac{\int_{S_L}^{S_{s_i}} [F(x)]^{m-1} dx}{[F(S_{s_i})]^{m-1}} \quad (5)$$

where m is the number of bidders. Recall that O_{s_i} is the same as U_b . From Equations 3, 4 and 5, the equilibrium bidding function for the seller can then be derived as follows:

$$P_{s_i}^* = C_{s_i} + \frac{\int_{S_L}^{S_{s_i}} [F(x)]^{m-1} dx}{[F(S_{s_i})]^{m-1}} \quad (6)$$

The seller in our mechanism also reasons about the expected future gain from winning the current auction. It takes into account the reputation of buyer b . In our mechanism, each buyer in the marketplace has a fixed number of neighbors that the buyer trusts and from which it can ask advice about sellers. This forms a social network of buyers where there is a directed link (edge) from a buyer to its neighbors. The edges are assigned weights $\in (0, 1]$ representing how much a buyer trusts its neighbors modeled using the approach that will be described in Section 3.2. A buyer becomes reputable if based on its reports about sellers it is accepted into the neighborhood of many other buyers. Cooperating with reputable buyers will allow the seller to build its own reputation and to be known as a trustworthy seller by many buyers in the marketplace. It will then be able to obtain more opportunities for doing business with buyers and to gain more profit in the future. We next provide formulae for the seller's reasoning about its expected future gain and prove that the expected future gain the seller s_i can earn after doing good business increases with the reputation of buyer b .

We define the global reputation of buyer b (denoted as $R_b \in (0, 1]$) on the social network to be the network effect of the buyer, which represents how much this buyer influences other buyers' decisions on the entire network. According to [5], reputation should be calculated as the effect that this buyer has on other buyers it influences, multiplied by these other buyers' effect on the network. This is a recursive and fixed-point computation. We use the following formula to compute the reputation of the buyer:

$$\bar{R} = \bar{L}^T \cdot \bar{R} \quad (7)$$

where \bar{R} is a vector containing each buyer's reputation. \bar{L} is an asymmetric matrix of the normalized weights of edges between all two-buyer pairs. The weight between two buyers is 0 if there is no link between them. Therefore, \bar{R} can be computed as the dominant eigenvector of \bar{L} .²

If the seller cooperates with the buyer, the new satisfied encounter between the buyer and the seller will then increase the seller's trustworthiness. The seller's

² \bar{R} can be recorded by the central server and shared with sellers.

probability of being allowed to join the buyer's auctions in the future will be increased by some amount, ΔP_b , where $\Delta P_b > 0$. Since this increment in probability is fairly small and relatively stable, we can assume that the probability of the seller being involved in auctions of other neighboring buyers increases linearly with how much these other buyers trust the current buyer b . The increase in probability of a seller being involved in every buyer's auctions across the network is $\Delta P_b R_b$.

If the seller is involved in a buyer's auction, the average probability of winning the auction is $\frac{1}{m}$, given that the number of bidders in the buyer's auction is m . The seller's average profit of being involved in a buyer's auction will then be $\frac{S_{s_i}}{m^2}$, which is the average probability of winning the auction multiplied by the average instant profit gained from winning the auction.³ We use $E_{s_i}(R_b)$ to denote the amount of the seller's expected future gain and $n' > 0$ the number of expected purchases from buyers. The expected future profit $E_{s_i}(R_b)$ is then

$$E_{s_i}(R_b) = n' \frac{S_{s_i}}{m^2} \Delta P_b R_b \quad (8)$$

From Equation 8, we have the following inequality:

$$\frac{\partial[E_{s_i}(R_b)]}{\partial R_b} = \frac{\partial[n' \frac{S_{s_i}}{m^2} \Delta P_b R_b]}{\partial R_b} = n' \frac{S_{s_i}}{m^2} \Delta P_b \geq 0 \quad (9)$$

The expected future gain the seller s_i can earn increases with the reputation of the buyer b .

Let us first consider a simplified scenario where sellers $\{s_1, s_2, \dots, s_m\}$ 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 following assumption that the distribution of S_{s_i} , $F(\cdot)$ is a uniform distribution. Let us also assume that the seller's lowest realized surplus S_L for a transaction is 0. Equation 6 can then be simplified as follows:

$$P_{s_i}^* = C_{s_i} + \frac{\int_{S_L}^{S_{s_i}} [F(x)]^{m-1} dx}{[F(S_{s_i})]^{m-1}} = C_{s_i} + \frac{\int_0^{S_{s_i}} (\frac{x}{S_H})^{m-1} dx}{(\frac{S_{s_i}}{S_H})^{m-1}} = C_{s_i} + \frac{S_{s_i}}{m} \quad (10)$$

Since the seller's realized surplus is equal to the sum of the buyer and the seller's profit and the seller has expected future gain from winning the current auction, the seller's realized surplus S'_{s_i} can then be changed as follows:

$$S'_{s_i} = U_b + U_{s_i} + \lambda E_{s_i}(R_b) = V_b - C_{s_i} + \lambda E_{s_i}(R_b) = S_{s_i} + \lambda E_{s_i}(R_b) \quad (11)$$

where $\lambda \in [0, 1]$ is a discounting factor.⁴ 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,

³ The average instant profit is $\frac{S_{s_i}}{m}$, as shown in Equation 10.

⁴ We suggest the inclusion of a discounting factor to allow sellers to learn over time the likelihood of receiving their expected future gain. The proofs that follow do not depend on its inclusion.

the symmetric Bayes-Nash equilibrium surplus offer function formalized in Equation 5 should be changed as follows:⁵

$$O_{s_i}^* = S_{s_i} + \lambda E_{s_i} - \frac{\int_{\lambda E_{s_i}}^{S'_{s_i}} [F(x)]^{m-1} dx}{[F(S'_{s_i})]^{m-1}} \quad (12)$$

From Equations 3, 4 and 12, we then can derive the modified equilibrium bidding function for the seller as follows:

$$\begin{aligned} P_{s_i}^* &= C_{s_i} - \lambda E_{s_i} + \frac{\int_{\lambda E_{s_i}}^{S'_{s_i}} [F(x)]^{m-1} dx}{[F(S'_{s_i})]^{m-1}} \\ &= C_{s_i} - \lambda E_{s_i} + \frac{\int_{\lambda E_{s_i}}^{S_{s_i} + \lambda E_{s_i}} \left(\frac{x}{S_H}\right)^{m-1} dx}{\left(\frac{S_{s_i} + \lambda E_{s_i}}{S_H}\right)^{m-1}} \\ &= C_{s_i} + \frac{S_{s_i}}{m} - \frac{1}{m} \left[\frac{(\lambda E_{s_i})^m}{(S_{s_i} + \lambda E_{s_i})^{m-1}} + (m-1)\lambda E_{s_i} \right] \end{aligned} \quad (13)$$

Comparing Equation 10 with Equation 13, we can see that the seller should offer the buyer reward $D_{s_i}(R_b)$ as follows:

$$D_{s_i}(R_b) = \frac{1}{m} \left[\frac{(\lambda E_{s_i})^m}{(S_{s_i} + \lambda E_{s_i})^{m-1}} + (m-1)\lambda E_{s_i} \right] \quad (14)$$

The reward can be the decreased price of the product.⁶ According to Equation 9, the seller's expected future gain $E_{s_i}(R_b)$ is a monotonically increasing function of R_b , the reputation of buyer b . We can then prove that the reward $D_{s_i}(R_b)$ offered to the buyer is also a monotonically increasing function of R_b , shown as follows:

$$\begin{aligned} \frac{\partial D_{s_i}}{\partial R_b} &= \frac{\partial \left\{ \frac{1}{m} \left[\frac{(\lambda E_{s_i})^m}{(S_{s_i} + \lambda E_{s_i})^{m-1}} + (m-1)\lambda E_{s_i} \right] \right\}}{\partial R_b} \\ &= \frac{1}{m} \left[\frac{\partial \left(\frac{(\lambda E_{s_i})^m}{(S_{s_i} + \lambda E_{s_i})^{m-1}} \right)}{\partial (\lambda E_{s_i})} \lambda \frac{\partial E_{s_i}}{\partial R_b} + (m-1)\lambda \frac{\partial E_{s_i}}{\partial R_b} \right] \\ &= \frac{\lambda}{m} \left[\frac{m(\lambda E_{s_i})^{m-1}}{(S_{s_i} + \lambda E_{s_i})^{m-1}} - \frac{(m-1)(\lambda E_{s_i})^m}{(S_{s_i} + \lambda E_{s_i})^m} + m-1 \right] \frac{\partial E_{s_i}}{\partial R_b} \\ &\approx \left\{ \frac{m(\lambda E_{s_i})^{m-1}}{(S_{s_i} + \lambda E_{s_i})^{m-1}} + (m-1) \left[1 - \left(\frac{\lambda E_{s_i}}{S_{s_i} + \lambda E_{s_i}} \right)^m \right] \right\} \frac{\partial E_{s_i}}{\partial R_b} > 0 \end{aligned} \quad (15)$$

We have now proved the following proposition:

Proposition 1 *Sellers are better off providing better rewards to reputable buyers in the case where all sellers have the same productivity.*

⁵ We replace $E_{s_i}(R_b)$ by E_{s_i} for a more concise formulation.

⁶ According to Equation 3, if the bidding price is fixed, the reward can also be the increased values of the product offered to the buyer.

The above analysis depends on the simplified assumption that sellers have the same productivity. We can generalize this result by removing this assumption. In this case, sellers may have different costs for producing the product with the same value of V_b . We first modify the seller's original equilibrium bidding function formalized in Equation 6 based on Equation 4, shown as follows:

$$P_{s_i}^* = V_b - S_{s_i} + \frac{\int_{S_L}^{S_{s_i}} [F(x)]^{m-1} dx}{[F(S_{s_i})]^{m-1}} \quad (16)$$

We then prove that the seller's original equilibrium bidding function is a monotonically decreasing function of S_{s_i} :

$$\begin{aligned} \frac{\partial P_{s_i}^*}{\partial S_{s_i}} &= \frac{\partial \left\{ V_b - S_{s_i} + \frac{\int_{S_L}^{S_{s_i}} [F(x)]^{m-1} dx}{[F(S_{s_i})]^{m-1}} \right\}}{\partial S_{s_i}} \\ &= \frac{\frac{\partial [\int_0^{S_{s_i}} F(x)^{m-1} dx]}{\partial S_{s_i}}}{[F(S_{s_i})]^{m-1}} - \frac{\frac{\partial [F(S_{s_i})]^{m-1}}{\partial S_{s_i}} \int_0^{S_{s_i}} [F(x)]^{m-1} dx}{[F(S_{s_i})]^{2m-2}} - 1 \\ &= 1 - \frac{(m-1) \frac{\partial F(S_{s_i})}{\partial S_{s_i}} [F(S_{s_i})]^{m-2} \int_0^{S_{s_i}} [F(x)]^{m-1} dx}{[F(S_{s_i})]^{2m-2}} - 1 \\ &= - \frac{(m-1) \frac{\partial F(S_{s_i})}{\partial S_{s_i}}}{[F(S_{s_i})]^m} \int_0^{S_{s_i}} [F(x)]^{m-1} dx < 0 \end{aligned} \quad (17)$$

Based on Equation 9, we can see that the seller's modified realized surplus S'_{s_i} formalized in Equation 11 will also increase as R_b increases:

$$\frac{\partial S'_{s_i}}{\partial R_b} = \frac{\partial [S_{s_i} + \lambda E_{s_i}(R_b)]}{\partial R_b} = \lambda \frac{\partial [E_{s_i}(R_b)]}{\partial R_b} > 0 \quad (18)$$

Therefore, the following proposition holds:

Proposition 2 *The seller's equilibrium bidding function is a monotonically decreasing function of R_b , which indicates that the seller will give more reward $D_{s_i}(R_b)$ to the buyers that are considered as more reputable in the marketplace.*

Buyer Honesty Here we prove the following proposition:

Proposition 3 *The seller strategy creates incentives for buyers to truthfully report the results of their business with sellers in order to become more reputable in the marketplace.*

From Equation 3, we first formalize the total profit gained by the buyer b from l times of doing business with sellers, shown as follows:

$$T_b = \sum_{k=1}^l U_{b,k} = \sum_{k=1}^l (V_{b,k} - P_{s_k}^*) \quad (19)$$

Based on Proposition 2 that a seller's equilibrium bidding function $P_{s_k}^*$ is a monotonically decreasing function of R_b , we then can prove that the buyer's total profit T_b will increase with the increase of its reputation R_b , as follows:

$$\frac{\partial T_b}{\partial R_b} = \frac{\partial[\sum_{k=1}^l (V_{b,k} - P_{s_k}^*)]}{\partial R_b} = \sum_{k=1}^l \frac{\partial V_{b,k}}{\partial R_b} - \sum_{k=1}^l \frac{\partial P_{s_k}^*}{\partial R_b} = - \sum_{k=1}^l \frac{\partial P_{s_k}^*}{\partial R_b} > 0 \quad (20)$$

since $\frac{\partial P_{s_k}^*}{\partial R_b}$ is negative (and considering $V_{b,k}$ as independent of R_b). Therefore, in order to gain more total profit, it is better off for the buyer to maintain high reputation. This can be achieved by participating in the social network and honestly reporting the results of its business with sellers. The value of honest reporting is demonstrated as well in the experiments of Figures 1(a) and 4(a) in Section 4.

3.2 Buyer Strategy to Promote Seller Honesty

In this section, we present an effective strategy for buyers to choose their business partners. Buyers using this strategy are able to gain more profit, which is further validated by experimental results presented in Section 3.2. We also discuss how this strategy creates incentives for sellers to deliver what they promised in bids.

Buyer Strategy To avoid doing business with possibly dishonest sellers, the buyer b in our mechanism first models the trustworthiness of sellers. Different existing approaches for modeling sellers' trustworthiness can be used here, for example the approach advocated by Zhang and Cohen [6] and the TRAVOS model proposed by Teacy et al. [7]. Both approaches propose to take into account the buyer's personal experience with the sellers as well as ratings of the sellers provided by other buyers. The buyer in our mechanism will allow only a number of the most trustworthy sellers to join the auction. Sellers about which the buyer b does not have information will also be allowed to join the auction with a small probability.

However, buyers may provide unfair ratings of sellers. Our mechanism allows the central server to maintain a fixed number of neighbors for each buyer: a list of the most trustworthy other buyers to this buyer, used to provide advice about sellers, in order to form a social network of buyers.⁷ The trustworthiness of these other buyers then also needs to be modeled for periodically updating the buyer's neighbor list. In the experiments presented in Section 4, the approach of Zhang and Cohen [6] is used for this purpose. This approach allows the central server to model how trustworthy a buyer privately feels another buyer (advisor) to be based on their ratings (0 or 1) for commonly rated sellers (where, briefly, an advisor is trustworthy if its ratings of sellers within limited time windows agree with those of the buyer). When the buyer has limited private knowledge of

⁷ Note for a new buyer, the central server randomly assigns to it some other buyers as its neighbors.

the advisor, the public reputation of the advisor will also be considered, based on all ratings for the sellers ever rated by the advisor held in the central server (where, briefly, an advisor is trustworthy if it consistently agrees with the ratings provided for sellers by others). Finally, the trustworthiness of the advisor will be modeled by combining the private and public values. The ones with the highest trustworthiness will be chosen by the central server to be the buyer's neighborhood; untrustworthy ones will therefore decrease their opportunities for rewards from sellers.

Seller Honesty Our idea of allowing the buyer to limit the number of selected bidders in its auctions is supported by Kim's results demonstrated in [8]. Kim claims that public tendering could lead to quality reduction by bidders; in contrast, selective tendering depending on bidders' trustworthiness may avoid such difficulties. More specifically, Kim proves that by using a buyer strategy as described above (modeling the trustworthiness of sellers and limiting the number of sellers that are considered), dishonest sellers will not be able to gain more total profit than that gained by honest sellers. Suppose that a dishonest winning seller s decides not to deliver its promise in its bid submitted to the buyer b in the current auction. Also suppose that the seller's equilibrium bidding price is P_s and C_s is the cost for s to produce the delivered product. By assuming that a dishonest seller will lose the chance to do business with the buyer in the future, the total profit gained by the seller s can then be formalized based on Equation 2:

$$T_s = U_s = P_s - C_s \quad (21)$$

The study of [8] does not consider the case where buyers form a social network. The seller therefore does not take into account the future profit gained by doing business with other buyers influenced by the feedback about the seller provided by the buyer b . In our case, the seller bids to sell the product to the buyer by also taking into account the future gain obtained by doing business with other buyers that consider b as their neighbor. The seller's expected gain is then greater than or equal to that in their case. Greater expected future gain leads to a larger realized surplus (see Equation 11). Based on the argument supported by Equation 17 that the seller's equilibrium bidding function is a monotonically decreasing function of its realized surplus, the seller's equilibrium bidding price P'_s should then be less than or equal to P_s . The profit that the seller s is able to earn will be less than or equal to the profit that it can earn in the case where sellers do not take into account the expected future gain obtained from other buyers in the marketplace:

$$T'_s = U'_s = P'_s - C_s \leq P_s - C_s = T_s \quad (22)$$

Honest sellers in both cases (taking future gain into account, or not) instead are able to gain the same amount of profit. The sellers in our mechanism decrease their instant profit, which will be complemented by their expected future gain. Based on the above analysis, honest sellers in our mechanism therefore will be

able to gain more total profit than that gained by dishonest sellers. Rational sellers desire profit and therefore will be honest. In conclusion, we have now proved the following:

Proposition 4 *The buyer strategy is able to promote seller honesty.*

4 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; buyers are more profitable when they participate, by providing ratings to others; limiting the number of bidders by buyers promotes seller honesty; buyers derive better profit when they use the ratings of sellers provided by neighbors.

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. As discussed in Section 3.2, buyers' neighbor lists will be updated.

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.

4.1 The Value of 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 computed using Equation 7. The results⁸ are shown in Figure 1(a). From this figure, we can see that the buyers providing the smaller percentages of unfair ratings 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 10 days when our marketplace converges, the

⁸ All experimental results in Section 4 are averaged over 500 rounds of the simulation.

changes of buyers' reputation 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 3. From Figure 1(b), we can see that buyers providing fewer unfair ratings gain more total profit. Our mechanism promotes buyer honesty (Proposition 3). Note that the profit difference of different types of buyers is fairly small. This is because buyers have at most 30 requests in total.




Fig. 1. (a) Reputation of Buyers being Honest vs. Dishonest; (b) Profit of Buyers being Honest vs. Dishonest; (c) Average Trust of Sellers being Honest vs. Dishonest

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 1(c), results indicate that sellers being dishonest more often 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. This is because they do not have much chance to do business with buyers and do not have many ratings. A seller without any ratings will have a default trust value of 0.5. We also record the total profit gained by different sellers using our new incentive mechanism. Results are shown in Figure 2(a). From this figure, we can see that sellers being honest more often gain more profit. Therefore, our mechanism promotes seller honesty (Proposition 4). 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 for sellers to be considered trustworthy to be very high. The sellers lying 25% are not considered as trustworthy sellers, therefore have few occasions to be selected as business partners by buyers.

4.2 Seller Strategy

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 model reputation of buyers and offer better rewards to reputable buyers. Another group of sellers do not model reputation of buyers and ask

for the same price from different buyers. Sellers in each group lie in different percentages (0%, 25% and 75%) of their business with buyers. We measure the average trust values of sellers from each group. Results shown in Figure 2(b) indicate that sellers modeling reputation of buyers have higher average trust values. We also measure the total profit gained by different sellers. Results in Figure 2(c) indicate that sellers are better off to model reputation of buyers and adjust prices of products according to buyers' reputation. By satisfying reputable buyers, these sellers are able to attract more buyers to do business with, in order to gain more profit.

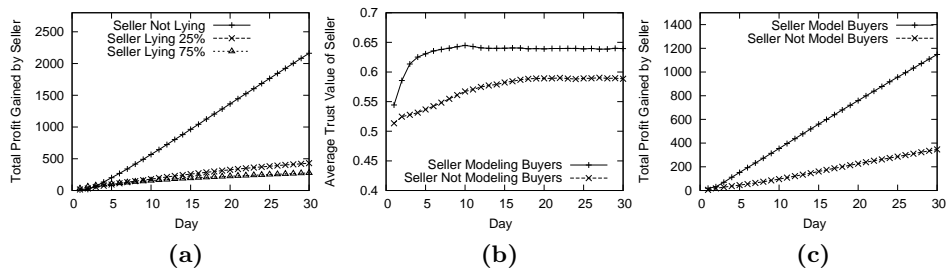


Fig. 2. (a) Total Profit Gained by Sellers being Honest vs. Dishonest; (b) Average Trust of Sellers Modeling vs. not Modeling Reputation of Buyers; (c) Total Profit Gained by Sellers Modeling vs. not Modeling Reputation of Buyers

4.3 Buyer Strategy

Buyers in the marketplace may also have different strategies. They may allow a lot of sellers to join their auctions. They may not always provide ratings for sellers. They may use different methods to model sellers, or may not model sellers at all. In this section, we carry out experiments to show the value of limiting the number of bidders and compare reputation values and total profit of buyers using different strategies. Results show that limiting the number of bidders promotes seller honesty, our mechanism provides incentives for buyers to provide ratings of sellers, and the modeling methods we propose provide buyers with more profit.

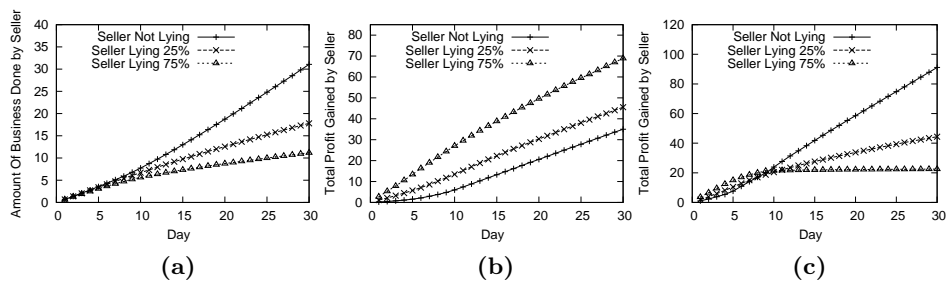


Fig. 3. (a) Sellers' Amount of Business When Allowing Many Sellers to Bid; (b) Total Profit Gained by Sellers When Allowing Many Sellers to Bid; (c) Total Profit Gained by Sellers When Limiting Seller Bids

Limiting Number of Bidders In this experiment, we have 90 sellers. Every 30 sellers acts dishonestly in different percentages (0%, 25% and 75%) of their business with buyers. In the first experiment, we allow 30 sellers to join each buyer’s auctions. Figure 3(a) shows the amount of business (number of transactions) done by different sellers. Sellers being honest more often are still able to gain more opportunities to do business with buyers. We also compare total profit gained by different sellers in this setting. However, from the results shown in Figure 3(b), we can see that sellers being dishonest more often gain more total profit. In this case, because more sellers are allowed to join buyers’ auctions, each seller’s equilibrium bidding price should be lower in order to win the auctions. Sellers being honest gain very little profit from each business with a buyer; therefore, dishonesty is promoted.

In the second experiment, we limit the number of bidders allowed in each of the buyers’ auctions to be 6. As shown in Figure 3(c), sellers being honest more often are able to gain more total profit. Honest sellers in this case are more likely to win the future auctions of buyers. They are offered sufficient future gain because limiting the number of bidders increases each seller’s equilibrium bidding price. Therefore, limiting the number of bidders promotes seller honesty.

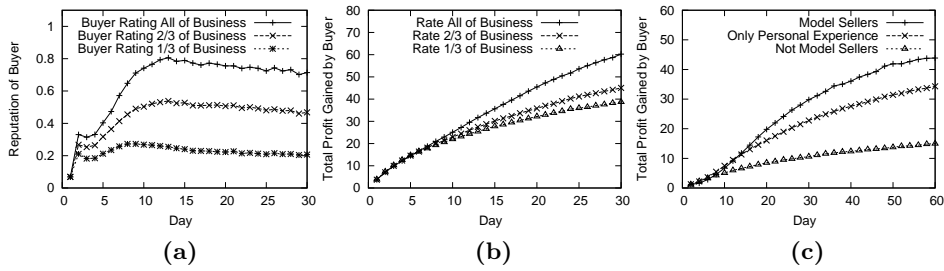


Fig. 4. (a) Reputation of Buyers Sharing Different Number of Ratings; (b) Total Profit Gained by Buyers Sharing Different Number of Ratings; (c) Total Profit Gained by Buyers Modeling vs. not Modeling Sellers

Incentives for Providing Ratings We examine how our mechanism 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 4(a). Buyers that provide more ratings have larger reputation values. We also measure total profit of these buyers. Results shown in Figure 4(b) indicate that buyers that have provided more ratings are able to gain more total profit. Therefore, it is better off for buyers to provide ratings of sellers. These ratings can then be shared with other buyers that consider these buyers as their neighbors, when modeling the trustworthiness 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 and advice about the sellers provided by their neighbors. Another third of the buyers uses only personal experience to model the trustworthiness of sellers. These buyers allow only a number of 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. We can see from Figure 4(c) that buyers modeling the trustworthiness of sellers and limiting their participation are 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 (an important element of our incentive mechanism) are 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.

5 Related Work

Side payment mechanisms [1, 9] have also been developed for promoting honesty in e-marketplaces. These mechanisms offer payment to buyers that fairly rate results of business with sellers. One facet of the side payment mechanisms in these papers is the requirement of a center to control monetary payments, so that budget balance is a concern. In contrast, in our mechanism the central server does not handle payments; rewards are directed from sellers to buyers. Brynov and Sandholm [10] design a mechanism that provides incentives for sellers to truthfully reveal their trustworthiness at the beginning of their business with buyers. Their mechanism has the limitation that the number of goods the buyers will purchase may not depend on their actual needs but has to be dependent on the trustworthiness of the seller.

The problem that strategic agents may collude with each other has been acknowledged as an important consideration by several researchers in the field (e.g [1]). Side payment mechanisms based simply on the similarity of buyers' ratings may therefore have difficulty with the situation where buyers collude in giving unfair ratings. Jurca and Faltings [11] investigate side payment mechanisms that can cope with collusion. 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, our mechanism's use of neighborhoods provides an avenue for excluding colluding buyers. Assuming that some honest buyers exist in the marketplace, they will use that honesty to gain rewards and will also serve to exclude the differing, colluding buyers from their neighborhoods. The use of neighborhoods also helps to detect and avoid dishonest, colluding sellers (those that differ from the predominant opinion).

6 Conclusions and Future Work

In this paper, we presented a detailed incentive mechanism to encourage honesty, intended for use in designing e-marketplaces. We provided theoretical proofs

to show that buyers have incentives to be honest in reporting about sellers, when sharing ratings with the buyers in their neighborhoods, under our particular framework. This occurs as a result of sellers offering better rewards to more reputable buyers, as part of their reasoning about how to obtain profit. We are also able to show that seller honesty is promoted, within our proposed framework, in order for sellers to receive higher profit. We further validated our mechanism through a set of experiments carried out using a simulated dynamic e-marketplace. As a result, our research emphasizes the value of using trust modeling and the sharing of reputation ratings in social networks in the design of an effective incentive mechanism.

In future work, we will carry out more extensive experimentation to continue to validate our model by comparing with others' models. In our future experiments, we will examine the situation where agents may vary their behavior widely to exploit the marketplace. In addition, we are particularly interested in empirically demonstrating how our framework is able to handle marketplaces where strategic agents collude with each other, more effectively than competing incentive-based trust models.

References

1. Jurca, R., Faltings, B.: An incentive compatible reputation mechanism. In: Proceedings of the IEEE Conference on E-Commerce. (2003)
2. Zhang, J., Cohen, R.: Design of a mechanism for promoting honesty in e-marketplaces. In: Proceedings of AAAI. (2007)
3. Gintis, H., Smith, E.A., Bowles, S.: Costly signaling and cooperation. *Journal of Theoretical Biology* **213** (2001) 103–19
4. Shachat, J., Swarthout, J.T.: Procurement auctions for differentiated goods. Experimental Economics Center Working Paper Series, Georgia State University (2006)
5. Richardson, M., Domingos, P.: Mining knowledge-sharing sites for viral marketing. In: Proceedings of KDD. (2002) 61–70
6. Zhang, J., Cohen, R.: A personalized approach to address unfair ratings in multi-agent reputation systems. In: Proceedings of the AAMAS Workshop on Trust in Agent Societies. (2006)
7. 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 Fourth International Autonomous Agents and Multiagent Systems. (2005)
8. Kim, I.G.: A model of selective tendering: Does bidding competition deter opportunism by contractors? *The Quarterly Review of Economics and Finance* **38**(4) (1998) 907–925
9. Miller, N., Resnick, P., Zeckhauser, R.: Eliciting informative feedback: The peer-prediction method. *Management Science* **51**(9) (2005) 1359–1373
10. Braynov, S., Sandholm, T.: Trust revelation in multiagent interaction. In: Proceedings of CHI'02 Workshop on the Philosophy and Design of Socially Adept Technologies. (2002) 57–60
11. Jurca, R., Faltings, B.: Collusion-resistant, incentive-compatible feedback payments. In: Proceedings of the Eighth ACM EC. (2007) 200–209