

COMBINING TRUST MODELING AND MECHANISM DESIGN FOR PROMOTING HONESTY IN E-MARKETPLACES

JIE ZHANG,¹ ROBIN COHEN,² AND KATE LARSON²

¹*School of Computer Engineering, Nanyang Technological University, Singapore*

²*School of Computer Science, University of Waterloo, Canada*

In this paper, we propose a novel incentive mechanism for promoting honesty in electronic marketplaces that is based on trust modeling. In our mechanism, buyers model other buyers and select the most trustworthy ones as their neighbors to form a social network which can be used to ask advice about sellers. In addition, however, sellers model the reputation of buyers based on the social network. Reputable buyers provide truthful ratings for sellers, and are likely to be neighbors of many other buyers. Sellers will provide more attractive products to reputable buyer to build their own reputation. We theoretically prove that a marketplace operating with our mechanism leads to greater profit both for honest buyers and honest sellers. We emphasize the value of our approach through a series of illustrative examples and in direct contrast to other frameworks for addressing agent trustworthiness. In all, we offer an effective approach for the design of e-marketplaces that is attractive to users, through its promotion of honesty.

Received 3 December 2010; Revised 6 April 2011; Accepted 9 May 2011; Published online 3 May 2012

Key words: buyer and seller honesty, electronic marketplace, trust and reputation, trust-based incentive mechanism.

1. INTRODUCTION

In electronic marketplaces that lack complete contracts and legal verification, a buying agent often relies on self-enforcing contracts where it can selectively choose business partners (selling agents) depending on their trustworthiness. This can create incentives for sellers to behave honestly in electronic marketplaces.

Modeling the trustworthiness of a seller can be based on the buyer's past personal experience with the seller (e.g., Tran and Cohen (2004)). However, for a new buyer or a buyer without much 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., Jøsang and Ismail (2002); Yu and Singh (2003); Teacy et al. (2005)). The problem of untruthful ratings may then arise. Buyers may provide untruthful high ratings to promote the seller. This is referred to as "ballot stuffing" (Dellarocas (2000)). Buyers may also provide untruthful low ratings, to cooperate with other sellers to drive a seller out of the marketplace. This is referred to as "bad-mouthing."

A variety of trust modeling approaches have been proposed to address the problem of untruthful ratings (Dellarocas (2000); Whitby, Jøsang, and Indulska (2005); Teacy et al. (2005); Zhang and Cohen (2006)). For example, the beta reputation system (BRS) of Whitby et al. (2005) propagates ratings provided by multiple buyers when estimating the reputation of a selling agent. It filters out those ratings that are not in the majority with other ratings. TRAVOS, developed by Teacy et al. (2005) proposes that possibly unreliable ratings of sellers need be discounted when the buying agent tries to reason about the trustworthiness of the sellers. While these methods for trust modeling can begin to address the problem of untruthful ratings, promoting honesty in the marketplace would be even more effective.

Address correspondence to Jie Zhang, School of Computer Engineering, Nanyang Technological University, Singapore; e-mail: zhangj@ntu.edu.sg

To promote buyer honesty, we propose a novel trust-based incentive mechanism. In our mechanism, buyers are encouraged to be truthful to gain more profitable transactions. This idea is supported by Gintis, Smith, and Bowles (2001). They argue that altruism in one context signals “quality” that is rewarded by increased opportunities in other contexts. Specifically, if the system is such that the provision of truthful reputation feedback makes agents more likely to choose to undertake transactions with the reporting agent, then the reporting agent would benefit for its feedback through a greater number of profitable transactions.

In our mechanism, we also allow sellers to explicitly model the reputability of buyers, based on the neighborhoods to which they belong in the society. A buyer is reputable if it is the neighbor of many other reputable buyers. Buyers that always provide truthful ratings of sellers are likely to become reputable. This is also supported by Gintis et al. (2001) through the model of a multiplayer 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”). Sellers in our system will increase quality and decrease prices of products to satisfy reputable buyers. This therefore creates an incentive for buyers to provide truthful ratings of sellers. Because buyers are sharing ratings of sellers, sellers are also encouraged to be trustworthy and honest (delivering the goods, as promised, to the buyers).

We assume a marketplace where buyers declare their interest in a good, sellers submit bids, and buyers ultimately select a seller with which to do business. We develop a precise formulation for sellers to reason about the important element of expected future profit, starting from formulae for reasoning about immediate profit. This is based on reasoning about the likelihood of making a sale to a buyer, making use of information held centrally about the reputation of the buyer. As a result, we are able to provide a precise 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. The proposed seller strategy and the buyer behavior in the context of the seller strategy are also illustrated through a detailed example in Section 4.

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. Thus, our mechanism is able to create a better environment for buyers and sellers to do business with each other, where honesty is promoted.

The remaining paper is organized as follows. We first present an overview of our proposed system, clearly described in terms of a diagram as well as buyer and seller algorithms. We then formalize our mechanism and provide theoretical proofs for buyer and seller honesty. We also introduce some examples. From here, we present our experimental results in simulated electronic marketplace environments. After that, we summarize existing incentive mechanisms for eliciting truthful ratings from buying agents and reflect on the advantages of our proposed mechanism in comparison with other competing approaches. Finally, we present conclusions and future work.

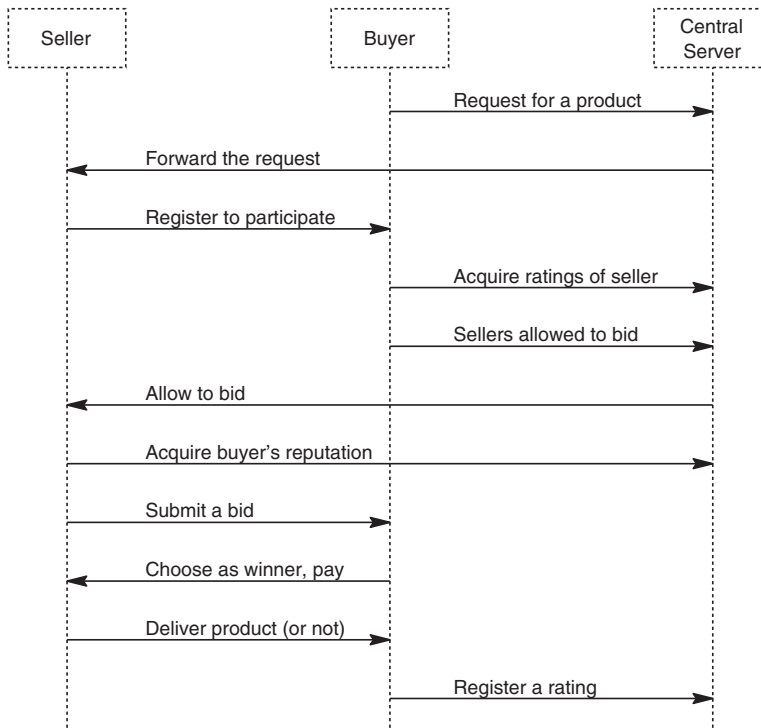


FIGURE 1. Buying and selling processes.

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. This server holds ratings of sellers submitted by buyers that will be shared with other buyers. It also forms a social network of buyers based on an approach that will be introduced in Section 5.3. The information about buyers' reputation is then also kept on the central server and will be released to sellers in the market.

Figure 1 illustrates the buying and selling processes, and the communications between buyers, sellers, and the central server. 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 that the buyer is interested in 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 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.

Algorithm 1. Buying Algorithm

```

1: Send a request for a product to the central server; //containing evaluation criteria for the product
2: Receive from the central server a list of sellers  $S$  interested in selling;
3: Set the list of sellers allowed to bid  $S' = \emptyset$ 
4: for  $s$  in  $S$  do
5:   Acquire ratings of  $s$  provided by neighbors from the central server;
6:   Model trustworthiness of  $s$ ; //using approaches that will be presented in Section 5.3
7:   if  $s$  is trustworthy then
8:     Add  $s$  in  $S'$ 
9:   end if
10: end for
11: Receive bids from each seller in  $S'$ ;
12: Choose the winner  $s_w$  that offers largest profit, pay to  $s_w$ ;
13: if  $s_w$  delivers promise then
14:   Submit a rating 1 to the central server;
15: else
16:   Submit 0;
17: end if

```

Sellers that are 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 buyer the largest profit, based on the buyer's evaluation criteria.

To formulate their bids, we introduce the important element that 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). As such, we are critically leveraging the social network of the buyers as part of the framework. As will be shown later in the next section, this eventually provides important incentives for honest reporting about sellers from 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). A pseudocode summary of the buying and selling algorithms is shown in Algorithms 1 and 2, respectively.

Algorithm 2. Selling Algorithm

```

1: Receive buyer  $b$ 's request from the central server;
2: if interested then
3:   Register to participate in  $b$ 's auction;
4: end if
5: if allowed to bid then
6:   Acquire reputation of  $b$  from the central server;
   //The central server calculates buyer reputation based on social network of buyers
   //The central server maintains neighbor lists
7:   Formulate a bid, and submit;
8: end if
9: if decide to be honest then
10:  Deliver the product described in the bid;
11: end if

```

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 from winning a buyer's auction. Next, we 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 lift 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.

3.1.1. Seller Strategy. We discuss our mechanism in the context of the request for quote (RFQ) system (Shachat and Swarthout (2006); Wan and Beil (2008)). We consider a scenario where a buyer b wants to buy a product p . The buyer specifies its evaluation criteria for a set of nonprice 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 nonprice feature. Each weight represents how much its corresponding nonprice feature is worth. A higher weight for a nonprice feature implies that the buyer cares more about the feature. The buyer also provides information in its evaluation criteria about the conversion from descriptive nonprice feature values to numeric values (for example, a 3-year warranty is converted to the numeric value of 10 on a scale of 1–10).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, which can be formalized as

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

follows:

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

We now begin to express 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 kinds of bids they should offer to buyers.

A seller s_i ($1 \leq i \leq m$) sets the price and values for the nonprice features of the product p (i.e., its promise), 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 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 nonprice 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:

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

The buyer's profit is also called the seller's "surplus offer," denoted as O_{s_i} .

The seller's "realized surplus" (Shachat and Swarthout 2006) is typically calculated as the sum of the buyer's and the seller's profit, as follows:

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

Note that the seller's realized surplus is higher when its cost for producing the product is lower. We also define the cumulative distribution function for the random variable $S_{s_i} > 0$ (over all sellers) as $F() \in [0, 1]$ and the support of $F()$ is $[S_L, S_H]$, where S_L and S_H are the lowest and highest possible values of the realized surplus of sellers, respectively. 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 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 by Shachat and Swarthout (2006), 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)–(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




FIGURE 2. A simple example of buyer social network.

which it can ask advice about sellers.² 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 approaches that will be described in Section 3.2. A simple example of buyer social network is shown in Figure 2. As shown in the figure, an arrow from b_3 to b_1 with the weight of 0.8 represents that buyer b_1 is b_3 's neighbor and the trust value b_3 has of b_1 is 0.8.

A buyer is reputable in the social network if it is the neighbor of many other reputable buyers. For example, buyer b_4 in Figure 2 is more reputable because it is highly trusted by buyer b_2 and the reputable buyer b_1 . 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 with b increases with the reputation of buyer b .

We define the global reputation of buyer b (denoted as R_b) 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 Richardson and Domingos (2002), reputation can 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 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 and is initially set to 1 for every buyer.³ An example of this calculation can be found in Section 4. \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} , which is similar to the EigenTrust computation (Kamvar, Schlosser, and Garcia-Molina 2003).⁴

If the seller cooperates with the buyer, the new satisfied encounter between the buyer and the seller will then increase the buyer's trust in the seller. The seller's probability of being allowed to join the buyer's auctions in the future will be increased by some amount, ΔP_b ,

² We require a fixed number of neighbors as part of our incentive mechanism.

³ \bar{R} can be recorded by the central server and shared with sellers.

⁴ Note that this calculation may be computationally intensive. A simpler way of calculating buyers' reputation can be to represent the reputation of a buyer based on the number of other buyers considering this buyer as one of their neighbors.

where $\Delta P_b > 0$. Because this increment in probability is small and 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 . Richardson and Domingos (2002) stated that 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 probability of winning the auction is estimated to be $\frac{1}{m}$, 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.⁵ We use $E_{s_i}(R_b)$ to denote the amount of the seller's expected future gain. The expected future profit $E_{s_i}(R_b)$ is then

$$E_{s_i}(R_b) = \frac{S_{s_i}}{m^2} \Delta P_b R_b. \tag{8}$$

From equation (8) and by assuming that $E_{s_i}(R_b)$ is differentiable, we have the following inequality:

$$\frac{\partial [E_{s_i}(R_b)]}{\partial R_b} = \frac{\partial \left[\frac{S_{s_i}}{m^2} \Delta P_b R_b \right]}{\partial R_b} = \frac{S_{s_i}}{m^2} \Delta P_b > 0. \tag{9}$$

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 S_{s_i} is the same for the sellers. The cumulative distribution function of S_{s_i} , $F()$ is then a linear function of S_{s_i} . Let us also assume that the seller's lowest realized surplus S_L for a transaction is close to 0. Equation (6) then can be simplified as follows:

$$\begin{aligned} P_{s_i}^* &= C_{s_i} + \frac{\int_{V_L - C_H}^{S_{s_i}} [F(x)]^{m-1} dx}{[F(S_{s_i})]^{m-1}} = C_{s_i} + \frac{\int_0^{S_{s_i}} \left(\frac{x}{S_H}\right)^{m-1} dx}{\left(\frac{S_{s_i}}{S_H}\right)^{m-1}} \\ &= C_{s_i} + \frac{x^m}{m(S_H)^{m-1}} \Big|_0^{S_{s_i}} = C_{s_i} + \frac{(S_{s_i})^m - 0}{m(S_{s_i})^{m-1}} = C_{s_i} + \frac{S_{s_i}}{m}. \end{aligned} \tag{10}$$

Because 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:

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

⁵ The average instant profit is $\frac{S_{s_i}}{m}$, as shown in equation (10).

where $\lambda' \in [0, 1]$ is a discounting factor.⁶ 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 (5) should be changed as follows:⁷

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

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

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

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

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

The reward can be the decreased price of the product. 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. 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

⁶ 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.

⁷ We replace $E_{s_i}(R_b)$ by E_{s_i} for a more concise formulation.

follows:

$$\begin{aligned}
\frac{\partial D_{s_i}}{\partial R_b} &= \frac{\partial \left\{ \frac{1}{m} \left[\frac{(\lambda' E_{s_i})^m}{(S_{s_i} + \lambda' E_{s_i})^{m-1}} + (m-1)\lambda' E_{s_i} \right] \right\}}{\partial R_b} \\
&= \frac{1}{m} \left[\frac{\partial \frac{(\lambda' E_{s_i})^m}{(S_{s_i} + \lambda' E_{s_i})^{m-1}}}{\partial (\lambda' E_{s_i})} \lambda' \frac{\partial E_{s_i}}{\partial R_b} + (m-1)\lambda' \frac{\partial E_{s_i}}{\partial R_b} \right] \\
&= \frac{\lambda'}{m} \left[\frac{\partial \frac{(\lambda' E_{s_i})^m}{(S_{s_i} + \lambda' E_{s_i})^{m-1}}}{\partial (\lambda' E_{s_i})} + (m-1) \right] \frac{\partial E_{s_i}}{\partial R_b} \\
&= \frac{\lambda'}{m} \left[\frac{m(\lambda' E_{s_i})^{m-1}(S_{s_i} + \lambda' E_{s_i})^{m-1}}{(S_{s_i} + \lambda' E_{s_i})^{2m-2}} \right. \\
&\quad \left. + m-1 - \frac{(m-1)(S_{s_i} + \lambda' E_{s_i})^{m-2}(\lambda' E_{s_i})^m}{(S_{s_i} + \lambda' E_{s_i})^{2m-2}} \right] \frac{\partial E_{s_i}}{\partial R_b} \\
&= \frac{\lambda'}{m} \left[\frac{m(\lambda' E_{s_i})^{m-1}}{(S_{s_i} + \lambda' E_{s_i})^{m-1}} - \frac{(m-1)(\lambda' E_{s_i})^m}{(S_{s_i} + \lambda' E_{s_i})^m} + m-1 \right] \frac{\partial E_{s_i}}{\partial R_b} \\
&\approx \left\{ \underbrace{\frac{m(\lambda' E_{s_i})^{m-1}}{(S_{s_i} + \lambda' E_{s_i})^{m-1}}}_{\geq 0} + (m-1) \underbrace{\left[1 - \left(\frac{\lambda' E_{s_i}}{S_{s_i} + \lambda' E_{s_i}} \right)^m \right]}_{> 0} \right\} \frac{\partial E_{s_i}}{\partial R_b} \quad (15) \\
&> 0.
\end{aligned}$$

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.

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 in equation (6) is a monotonically decreasing function of S_{s_i} where $\frac{\partial F(S_{s_i})}{\partial S_{s_i}} > 0$ because $S_{s_i} > 0$:

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

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 more reputable in the marketplace.

3.1.2. 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 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:

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

because $\frac{\partial P_{s_k}^*}{\partial R_b}$ is negative (and considering $V_{b,k}$ as independent of R_b). Therefore, 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.

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 5. We also discuss how this strategy creates incentives for sellers to deliver what they promised in their bids.

3.2.1. 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 (2006) and the TRAVOS model proposed by Teacy et al. (2005). 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.

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

However, buyers may provide untruthful ratings of sellers. Our mechanism allows the central server to maintain a fixed number⁸ of neighbors for each buyer: a list of the most trustworthy other buyers to this buyer, used to provide advice about sellers, to form a social network of buyers. The trustworthiness of these other buyers (advisors) then needs to be modeled. In the experiments presented in Section 5, the approach of Zhang and Cohen (2006) is used for this purpose. This approach allows a buyer to first model private reputation of an advisor based on their ratings 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 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 reputation values.

⁸ Deciding the appropriate number of neighbors to use is left for future work.

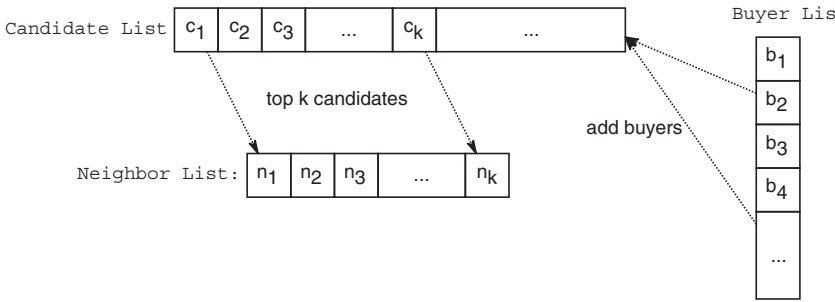


FIGURE 3. Candidate list and neighbor list.

Figure 3 shows an example of a candidate list and a neighbor list for a particular buyer. Assume that these are both ordered from most trustworthy to least trustworthy. For a new buyer, the central server randomly assigns to it some other buyers with high public reputation as candidates for its neighbors. The candidate list is larger than the neighbor list.⁹ The neighbor list will be updated periodically. Each time, the most trustworthy candidates will be selected as neighbors (shown as arrows from the candidate list to the neighbor list in Figure 3). 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 (shown as arrows from the buyer list to the candidate list in Figure 3).

3.2.2. 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 Kim (1998). Kim proves that public tendering could lead to quality reduction by bidders; in contrast, selective tendering depending on bidders’ trustworthiness may avoid such difficulties. Calzolari and Spagnolo (2006) also analyze repeated procurement processes. They show that by limiting the number of competitors and carefully choosing the trustworthy ones to join their auctions, buyers offer sellers sufficient future gain so that sellers will prefer to provide acceptable levels of quality of products in the current auction to build their reputation, to gain more profit in the future. Bar-Isaac also uses an example in Bar-Isaac (2005) to show that low competition may sustain an equilibrium in which sellers produce high-quality products.

In Kim (1998) and Calzolari and Spagnolo (2006), the authors prove 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 (possibly zero). By assuming that a dishonest seller will lose the chance to do business with the buyer in the future (i.e., no future profit), the total profit gained by the seller s can then be formalized based on equation (2), as follows:

$$T_s = U_s = P_s - C_s. \tag{21}$$

The studies of Kim (1998) and Calzolari and Spagnolo (2006) do 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

⁹ The suggestion of keeping track of a longer candidate list is also used in Yu and Singh (2000).

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 in our case 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:

$$T'_s = U'_s = P'_s - C_s \leq P_s - C_s = T_s. \tag{22}$$

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. EXAMPLES

In this section, we use some examples to demonstrate how our mechanism works. We first provide an example to demonstrate how a buyer selects the winning seller to do business with, based on not only the sellers' bids but also their trustworthiness. We then provide another example to illustrate how a seller models reputation of buyers and specifies its bids for buyers' requests according to their reputation values.

4.1. Buyer Choosing Winning Seller

In this example, a buyer b wants to buy a product p . It sends the request to the central server. In its request, the buyer specifies the two nonprice features (delivery time and warranty) of the product p ; the weight for each nonprice feature and the information about the conversion from descriptive nonprice feature values to numeric values are presented in Table 1. For example, the delivery time of 1 week will be converted to the numeric value of 3.

The central server forwards b 's request to the sellers in the marketplace. There are five sellers $\{s_1, s_2, s_3, s_4, s_5\}$ that are interested in selling their products to the buyer. The

TABLE 1. Buyer b 's Evaluation Criteria for p .

Nonprice 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

TABLE 2. Sellers Bidding for b 's Request.

Seller	Cost	S_s	P_s^*	O_s^*
s_3	5	5	6.06	3.94
s_4	6	4	6.72	3.28
s_5	8	2	8.04	1.96

trustworthiness of the sellers is assumed to be as follows:

$$Tr(s_1) = 0.39, \quad Tr(s_2) = 0.5, \quad Tr(s_3) = 0.83, \quad Tr(s_4) = 0.72, \quad Tr(s_5) = 0.72.$$

We set the threshold for sellers to be considered as trustworthy to be 0.7. In this case, only the sellers $s_3, s_4,$ and s_5 will be considered as trustworthy sellers by buyer b , and are allowed to submit their bids to the buyer. Suppose that all three sellers want to produce the same product for the buyer, which has 3 year warranty and will be delivered in 1 day. The buyer's value for their products will be calculated using equation (1) as follows:

$$V_b = 10 \times 0.4 + 10 \times 0.6 = 10.$$

The sellers $s_3, s_4,$ and s_5 have different costs for producing the product p . The realized surplus of each seller S_s calculated using equation (4), the sellers' equilibrium bidding price P_s^* calculated using equation (13), and their surplus offer for the buyer O_s^* calculated using equation (12) are listed in Table 2. In this example, we simplify the calculation by assuming that the sellers' expected future gain from winning the buyer's current auction is 1; we also set the discounting factor λ to 0.9. A detailed example is in Section 4.2 to show how a seller reasons about its expected future gain from winning the current auction.

The buyer b will choose the seller that has the largest surplus offer O_s as the winner of the auction. In this case, s_3 will be the winner. The buyer pays 6.06 to seller s_3 . Later on, seller s_3 delivers the product. Suppose that the seller delivers the product with 3 year warranty in 1 day; we say that the seller is trustworthy in this transaction. Buyer b will submit a rating of 1 to the central server. From this example, we can see that only the trustworthy seller s_3 gains the instant profit, which can be calculated according to equation (2) as follows:

$$U_{s_3} = P_{s_3} - C_{s_3} = 6.06 - 5 = 1.06.$$

4.2. Seller Bidding for Buyers' Requests

In this example, we illustrate how reputation of buyers is modeled by the central server and how a seller s specifies its bids for buyers' requests according to their reputation values. Suppose that there are six 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 1.

Assume that each buyer is allowed to have only three neighbors in this example. The neighbors of each buyer are listed in Table 3. We also assume that the trust value each buyer has of each its neighbor is 0.8. We calculate each buyer's reputation before normalization

TABLE 3. Neighbors of Buyers.

Buyer	Neighbors		
b_1	b_2	b_5	b_6
b_2	b_4	b_5	b_6
b_3	b_4	b_5	b_6
b_4	b_3	b_5	b_6
b_5	b_3	b_4	b_6
b_6	b_3	b_4	b_5

TABLE 4. Products Produced for Different Buyers.

Buyers	Nonprice 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

using equation (7) where

$$\bar{L} = \begin{bmatrix} & b_1 & b_2 & b_3 & b_4 & b_5 & b_6 \\ b_1 & - & 0.33 & 0 & 0 & 0.33 & 0.33 \\ b_2 & 0 & - & 0 & 0.33 & 0.33 & 0.33 \\ b_3 & 0 & 0 & - & 0.33 & 0.33 & 0.33 \\ b_4 & 0 & 0 & 0.33 & - & 0.33 & 0.33 \\ b_5 & 0 & 0 & 0.33 & 0.33 & - & 0.33 \\ b_6 & 0 & 0 & 0.33 & 0.33 & 0.33 & - \end{bmatrix} .$$

After one iteration of the calculation, the reputation of buyers becomes:

$$R_{b_1} = 0.1, \quad R_{b_2} = 0.39, \quad R_{b_3} = 1.0, \\ R_{b_4} = 1.3, \quad R_{b_5} = 1.6, \quad R_{b_6} = 1.6.$$

And after the calculation converges, the final reputation of buyers becomes

$$R_{b_1} = 0.8, \quad R_{b_2} = 1.31, \quad R_{b_3} = 10.95, \\ R_{b_4} = 11.46, \quad R_{b_5} = 11.74, \quad R_{b_6} = 11.74.$$

Seller s needs to decide how to bid for each buyer's request. It considers the reputation of each buyer. 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 4 lists the buyers' values for the products, calculated using equation (1) based on Table 1. The seller s has different costs for producing these products, which are also listed in Table 4.

TABLE 5. Seller’s Prices for Different Buyers.

Buyer	b_1	b_2	b_3	b_4	b_5	b_6
$E_s(R_b)$	0.04	0.06	0.49	0.51	0.52	0.52
S'_s	2.04	2.06	2.49	2.51	2.52	2.52
D_s	0.027	0.04	0.333	0.347	0.354	0.354
P_s^*	1.640	1.627	3.334	3.319	8.313	8.313

From Table 4, we can see that the seller’s realized surplus before considering its expected future profit is 2 for every buyer. Suppose that there are three sellers in each auction ($m = 3$). Also, assume that the increase in probability ΔP of the seller being involved in a buyer’s action in the future if the seller satisfies the buyer is 0.2. We can calculate the seller’s expected future profit from winning buyer b_1 ’s auction, according to equation (8) as follows:

$$E_s(R_{b_1}) = \frac{2}{3^2} * 0.2 * 0.8 = 0.04.$$

Table 5 lists the seller’s amount of expected future gain $E_s(R_b)$ from selling the products to each of the six buyers with different reputation values. We assume the discounting factor λ' to be 1. We also calculate the modified realized surplus S'_s using equation (11), and the reward D_s offered to different buyers and the seller’s equilibrium bidding prices P_s^* according to equation (13), as presented in Table 5.

We can see from Table 5 that seller s offers the best rewards to the more reputable buyers b_5 and b_6 . Buyers b_1 and b_2 with reputation values that are close to 0 gain very little reward. According to Tables 4 and 5, we can calculate the profit gained by the buyers using equation (3), as follows:

$$U_{b_1} = 1.360, \quad U_{b_2} = 1.373, \quad U_{b_3} = 1.666,$$

$$U_{b_4} = 1.681, \quad U_{b_5} = 1.687, \quad U_{b_6} = 1.687.$$

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, to gain more profit.

This example emphasizes the value of our proposed framework. When buyers are asked for their advice, they are inclined to be honest because failing to do so will impact the rewards that they ultimately receive from sellers. This is because the trustworthiness of each advisor is constantly being modeled, and untrustworthy advisors may then be dropped from neighborhoods. Because seller rewards are explicitly related to the number of neighborhoods in which the buyer has been accepted as an advisor, this ultimately affects the rewards that the advisor receives.

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

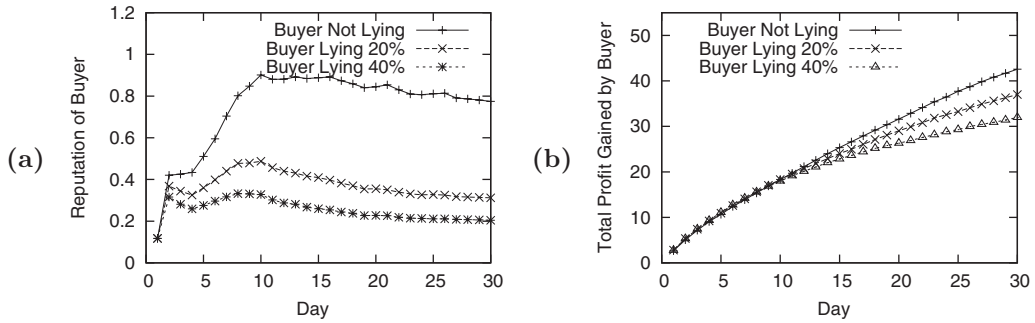


FIGURE 4. (a) Reputation of buyers being honest versus dishonest. (b) Profit of buyers being honest versus dishonest.

derive better profit when they use the ratings of sellers provided by neighbors and measure the trustworthiness of other buyers, 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 have 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 nonprice features. After they finish business with sellers, buyers rate sellers. Some buyers will provide untruthful ratings. Each group of buyers provides different percentages (0%, 20%, and 40%) of untruthful ratings. We allow two buyers from each group to leave the marketplace at the end of each day. Accordingly, we also allow six buyers to join the marketplace at the end of each day. These buyers will also provide a different percentage (0%, 20%, and 40%) of untruthful ratings, to keep the number of buyers in each group the same. Initially, we randomly assign five buyers to each buyer as its neighbors.

There are also nine sellers in total in the marketplace. Each three sellers act 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 nonprice 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 untruthful ratings. In our experiments, a buyer's reputation is computed using equation (7). The results¹⁰ are shown in Figure 4(a). From this figure, we can see that the buyers providing the smaller percentages of untruthful 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 10 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 untruthful ratings. The profit gained by a buyer from buying a product is formalized in

¹⁰ All experimental results in Section 5 are averaged over 500 rounds of the simulation.




FIGURE 5. (a) Average trust of sellers being honest versus dishonest. (b) Total profit of sellers being honest versus dishonest.

equation (3). From Figure 4(b), we can see that buyers providing fewer untruthful ratings will gain more total profit. This is consistent with the illustration outlined in Section 4.2 where buyers (such as b_1 and b_2) that were dishonest in their advice failed to be accepted by a sufficient number of other buyers, so belonged to fewer neighborhoods and thus received lesser rewards. Note that the profit difference of different types of buyers is small. This is because buyers have at most 30 requests in total. In summary, it is better off for buyers to provide truthful ratings of 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 5(a), 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 0.5.¹¹ 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 5(b). 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% 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% will not be considered as trustworthy sellers, therefore will have few occasions to be selected as business partners by buyers.

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.

¹¹ Note that 25% of the time these sellers are honest and do gain some trust.

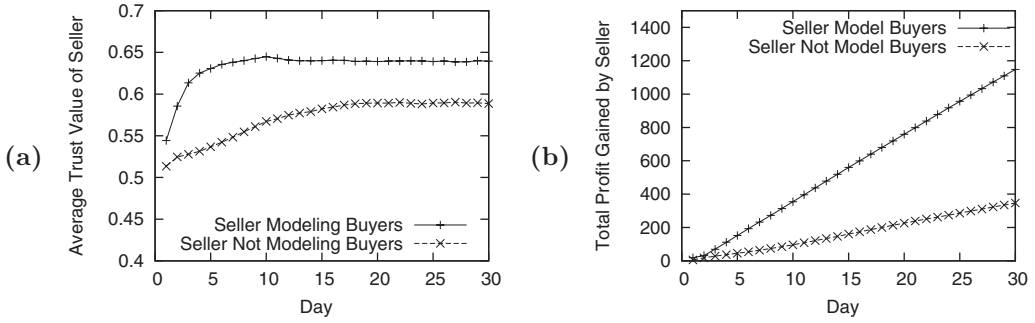


FIGURE 6. (a) Average trust value of different sellers. (b) Total profit gained by different sellers.

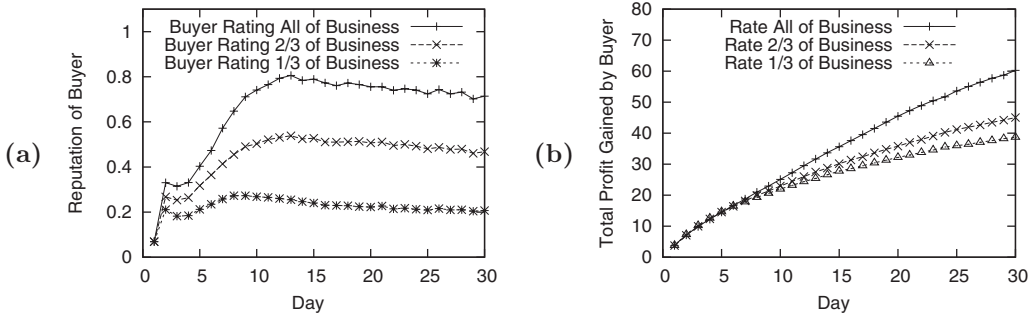


FIGURE 7. (a) Reputation of different buyers. (b) Profit gained by different buyers.

We measure the average trust values of sellers from each group. Results shown in Figure 6(a) indicate that sellers modeling reputation of buyers will have higher average trust values. We also measure the total profit gained by different buyers. Results in Figure 6(b) indicate that sellers are better off to model reputation of buyers and adjust prices of products according to buyers' reputation, 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 allow a lot of sellers to join their auctions. 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. Results show that our mechanism provides incentives for buyers to provide ratings of sellers, buyers should limit the number of bidders, and the modeling methods we propose will provide buyers with more profit.

5.3.1. 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 numbers of ratings. In this experiment, all buyers are honest. They have the same number of requests. However, they rate a different fraction (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(a). Buyers that have provided more ratings will have larger reputation values. We also measure total

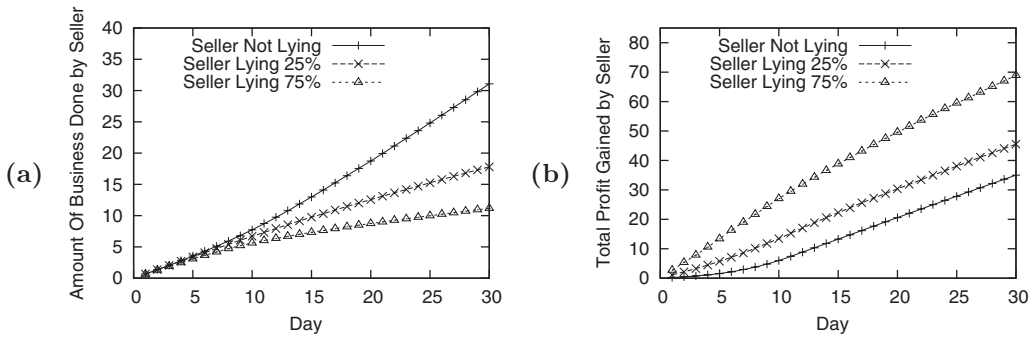


FIGURE 8. (a) Amount of business done by sellers. (b) Total profit gained by sellers.




FIGURE 9. Total profit gained by sellers.

profit of these buyers. Results shown in Figure 7(b) indicate that buyers that have provided more ratings will be able to gain more total profit because of lower prices offered by sellers. Therefore, it is better off for buyers to provide ratings of sellers.

5.3.2. Limiting Number of Bidders. We carry out experiments to show the importance of limiting seller bids. In these experiments, we have 90 sellers. Each 30 sellers act 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 8(a) shows the number of business 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 8(b), we can see that sellers being dishonest more often will gain more total profit. In this case, sellers being honest gain very little profit from each business with buyers; therefore, dishonesty will be promoted.

In the second experiment, we limit the number of bidders allowed in each of the buyer’s auctions to be 6. As shown in Figure 9, sellers being honest more often will be able to gain more total profit. Therefore, limiting the number of bidders allowed will promote seller honesty.

5.3.3. 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 on advice about the sellers provided by their neighbors. Another one-third of the buyers use




FIGURE 10. (a) Profit gained by different buyers. (b) Profit gained by different buyers.

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 remaining 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 10(a). From this figure, we can see that buyers modeling the trustworthiness of sellers and limiting their participation 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.

5.3.4. 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 untruthful ratings, but the second and third groups provide 20% and 40% of untruthful 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 untruthful 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(b). 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 untruthful 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. DISCUSSIONS

6.1. Related Work

Researchers have been developing incentive mechanisms to elicit truthful ratings from buyers. We survey three types of incentive mechanisms for electronic marketplaces, including

side payment mechanisms (Jurca and Faltings 2003; Miller, Resnick, and Zeckhauser 2005), credibility mechanisms (Jurca and Faltings 2004; Papaioannou and Stamoulis 2005), and trust revelation mechanisms (Braynov and Sandholm 2002; Dash, Ramchurn, and Jennings 2004). We also point out certain shortcomings of these methods, some of which motivated our work.

6.1.1. Side Payment Mechanism. Side payment mechanisms offer side payment to buyers that truthfully rate results of business with sellers. Providing truthful feedback of sellers is a Nash Equilibrium in these mechanisms.

In the incentive, compatible mechanism proposed by Jurca and Faltings (2003), a set of broker agents called R-agents, can sell and buy ratings of sellers to and from other ordinary agents. These ordinary agents first buy ratings from broker agents. After they finish doing business with sellers, they can sell ratings of the sellers back to the broker agents from which they bought ratings. An agent will get paid only if a rating of a seller it provides is the same as the next rating of the same seller provided by another agent. The price that an agent will get paid for a truthful report is determined based on the probability that an agent will trust another agent and this other agent will not change its behavior. However, they assume that broker agents already store some reputation information after bootstrapping the system. This overly simplifies the process of reputation management and additionally does not take into account the case of new entrants into the system. Moreover, this mechanism does not work if most of the agents provide untruthful feedback, an issue that may arise if agents are colluding.

Miller et al. (2005) introduce a mechanism which is similar to that proposed by Jurca and Faltings (2003). However, in this mechanism, 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 (Johnson, Pratt, and Zeckhauser 1990) 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 where every agent is better off providing truthful feedback, given that every agent else chooses the same strategy. This mechanism assumes that sellers have fixed quality, which limits its usefulness. As with the mechanism proposed by Jurca and Faltings (2003), the truthful equilibrium is not the only equilibrium in this mechanism. Because there may be nontruthful equilibria, where every agent is better off providing untruthful feedback, this mechanism is open once more to the challenge of buyers colluding to report untruthfully.

6.1.2. Credibility Mechanism. Instead of giving instant payment to agents that provide truthful ratings, credibility mechanisms measure agents' credibility or noncredibility 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 (2005) for eliciting truthful ratings in peer-to-peer systems. Besides reputation information, each peer also stores a noncredibility 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 noncredibility values will be decreased by the system. Otherwise, their noncredibility values will be increased and both of them will be punished. They will be forced not to conduct any transactions for a period that is exponential in their noncredibility values. The punishment of not transacting with other peers causes the punished peer to lose value offered by others. This provides incentives for peers to truthfully report the result of their business with others.

A slightly different credibility mechanism called “CONFESS” is proposed by Jurca and Faltings (2004) for the online hotel booking industry. This mechanism is used to cope with opportunistic behavior where a hotel may establish an excellent reputation first and then start cheating from time to time. It is based on the observation that hotels are less likely to cheat on clients that have a good reputation for reporting the truth, as the resulting negative report will attract future loss that outweighs the momentary gain obtained from cheating. One weakness of this mechanism is that it is difficult to guarantee budget balance among different involved parties. The center in this mechanism may end up earning a lot of extra profit, and the trustworthiness of the center becomes very crucial. The center may be incentivized to manipulate hotels and clients to gain more profit.

6.1.3. Trust Revelation Mechanism. Trust revelation mechanisms are designed to provide incentives for agents to truthfully report either the trustworthiness of themselves or trust values they place on other agents. These mechanisms are different from the ones of side payment and credibility. In side payment and credibility mechanisms, agents are asked to provide ratings of others or themselves that are binary (e.g., 1 or 0). Trust revelation mechanisms accept reported trust values that are continuous in the range of, for example, $[0, 1]$.

Braynov and Sandholm (2002) design a trust revelation mechanism that provides incentives for sellers to truthfully reveal their trustworthiness at the beginning of their business with buyers. This mechanism involves one trustworthy buyer and one possibly untrustworthy seller and operates as follows. At the beginning of business, the seller declares its trustworthiness. After that, the buyer chooses a quantity value in the business, for example, the quantity of the commodity that the buyer will purchase from the seller in this business transaction. This quantity value is dependent on the seller’s declared trustworthiness. If the quantity value is set properly, the seller will have the incentive to truthfully reveal its trustworthiness. This incentive mechanism has limited applicability, in that buyers have less control over the quantity of goods they want to purchase. In this case, the number of goods the buyers will purchase cannot depend on the buyers’ actual needs but has to be dependent on the trustworthiness of the seller.

Dash et al. (2004) also propose a trust revelation mechanism that explicitly handles issues of trust through mechanism design. The proposed mechanism is used in a task allocation scenario where agents need to make decisions about which other agents they should allocate their tasks to. It generalizes the standard Vickrey–Clarke–Groves (VCG) mechanism. In this mechanism, agents take into account the trustworthiness of other agents when determining their allocations. Each agent reports as well the trust that it places on other agents and its trust calculation function used to calculate the trustworthiness of another agent from all other agents’ reporting about the trustworthiness of the agent currently being modeled. The center will compute the trust of each agent according to this function and reported trust values of the agent from all other agents. The center then decides the optimal allocation and payment by using the intuition behind VCG mechanisms that an agent’s reporting of others’ trust affects only the allocation but not its payment to make sure that there are no incentives for agents to lie about their reporting of others’ trustworthiness. This approach assumes that all of an agent’s preferences concern its own allocation. However, buyers may provide inaccurate trust information to decrease or increase the chances of another agent receiving a good allocation, another scenario of collusion which poses a challenge.

6.1.4. Summary of the Related Work. We surveyed three different kinds of incentive mechanisms that are designed to provide incentives for agents to provide truthful reporting of trustworthiness, including side payment mechanisms, credibility mechanisms, and trust

revelation mechanisms. Side payment mechanisms offer payment to buyers that provide truthful ratings. Providing truthful ratings in these mechanisms is a Nash equilibrium. Credibility mechanisms measure agents' credibility. Agents in these mechanisms have incentives to provide truthful ratings, to increase their credibility or decrease their noncredibility. In doing so, they are able to gain higher profit. Trust revelation mechanisms have agents truthfully report their own trustworthiness or the trust they have of others that are represented as continuous values. In the trust revelation mechanism of Braynov and Sandholm (2002), selling agents are incentivized to truthfully report their own trustworthiness to obtain more business with buyers and gain more profit. And in the trust revelation mechanism of Dash et al. (2004), agents do not have incentives to lie about the trust they place on others.

One important issue that poses a challenge for all of these incentive mechanisms is collusion—scenarios where agents may be working together to provide unfair ratings. This can occur either when buyers collude with sellers or when buyers simply enter into collusive relationships with each other. While this problem is very difficult to combat, our model was motivated in part by being able to make some progress in coping with collusion, to a greater extent than competing approaches. In the following subsection, we move forward to contrast our model with those described in this section.

6.2. Advantages of Our Mechanism

In this subsection, we revisit our mechanism described in Section 3 and discuss some of its central advantages.

6.2.1. Wider Applicability. In our mechanism, a trustworthy seller is rewarded by a buyer from doing business with many more other buyers in the marketplace that consider this buyer as their neighbor. Therefore, unlike the trust revelation mechanism of Braynov and Sandholm (2002), our mechanism does not rely on the assumption that the quantity of the goods a buyer wants to purchase has to be dependent on sellers' trustworthiness. Different from some mechanisms, our mechanism also does not rely on the central server to handle monetary payments. Rewards are directed from sellers to buyers. It therefore easily achieves budget balance.

6.2.2. Addressing Collusion. The problem of coping with strategic agents that may collude with each other has been acknowledged as an important consideration by several researchers in the field (e.g., Jurca and Faltings 2003). Various elements of our particular incentive mechanism may provide an avenue for addressing collusion more effectively than other researchers.

Side payment mechanisms (Dellarocas 2002; Jurca and Faltings 2003; Miller et al. 2005) surveyed in Section 6.1.1 offer side payment to buyers that provide ratings of sellers that are similar to those provided by other buyers. These mechanisms have difficulty with the situation where buyers collude in giving untruthful ratings, to appear to be similar, to receive better rewards. Honest buyers are penalized because their truthful ratings are different from others' ratings. In contrast, in our mechanism, honest buyers will not be adversely affected by collusion in the marketplace; with our personalized approach for modeling the trustworthiness of advisors (Zhang and Cohen 2006), 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.

With further investigation, Jurca and Faltings (2007) observe that side payment mechanisms can cope with some collusion scenarios under some assumptions. For example, in the

scenario where all buying agents may collude but do not share revenues, their mechanism with a certain small amount of truthful feedback about sellers will be able to cope with such collusion. In the scenario where only some of the buyers collude but using different strategies and also distributing revenues, the side payment mechanism copes with collusion but only when the number of colluding buyers is small enough and other buyers are reporting honestly. In addition, all these scenarios only concern the case where buyers collude with each other. They do not consider the case where a seller may collude with a group of buyers in promoting the seller itself or in bad-mouthing another seller.

Credibility mechanisms (Papaioannou and Stamoulis 2005; Jurca and Faltings 2004) introduced in Section 6.1.2 suffer when buyers and sellers collude to increase each other's credibility. Because our mechanism allows the central server to maintain for buyers 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 should not impact any future decisions. In addition, all buyers have incentives to be honest, to enjoy the rewards offered by the honest sellers of the marketplace, if they maintain their position in many neighborhoods of the social network.

To explain how we begin to address collusion, consider the following simple example, where a buyer b needs to decide whether to do business with a seller s . The buyer models the trustworthiness of the seller. Suppose that the seller is, in fact, dishonest. A set of buyers $\{b_1, b_2, b_3, b_4\}$ colludes with seller s in promoting the seller by each providing the rating of 1. In this case, these buyers also collude with each other in providing the same ratings and will be rewarded by the side payment mechanisms. The seller will be rewarded by the credibility mechanisms because it will certainly provide a rating of 1 for itself. Another set of buyers $\{b_5, b_6\}$ honestly report the rating of 0 for the seller. Our mechanism makes use of the personalized approach for buyer b to model the trustworthiness of other buyers. Only the trustworthy buyers are considered by b as its neighbors from which it will ask advice about seller s . The dishonest buyers $\{b_1, b_2, b_3, b_4\}$ are subsequently excluded from the buyer's neighbor list. Therefore, their untruthful ratings will not mislead the buyer's decision about future sellers. If the dishonesty of b_1, b_2, b_3 , and b_4 is similarly detected by additional buyers, these untruthful advisors will continue to have low reputation and will not get high rewards from sellers. Note that the ratings provided by the honest buyers $\{b_5, b_6\}$ in the buyer's neighbor list will still allow the buyer to correctly model the trustworthiness of the seller and avoid doing business with the seller. The dishonest seller thus cannot gain better profit in our mechanism.

6.2.3. The Value of Trust Modeling. Our trust-based incentive mechanism is built on trust modeling to form a social network of buyers. The use of trust modeling presents some unique advantages for our incentive mechanism. An important assumption in the design of an incentive mechanism is that all agents are rational (Jurca and Faltings 2003). They all have the goal of maximizing their profit. However, there might be some agents that are irrational. In this case, our trust modeling approach becomes important. It can help agents model the trustworthiness of other agents and make correct decisions even when some irrational agents act dishonestly.

Another issue with incentive mechanisms is the exit problem, as acknowledged in Dellarocas (2002). If an agent plans to leave the market, it can cheat freely without repercussions (Kerr and Cohen 2007). In this case, this agent may not have incentives to be honest and incentive mechanisms may fail. However, if trust modeling is still used by buying agents,

it will be possible to detect dishonest advisors even if they plan to exit, to continue to make effective decisions about selling agents.

6.3. The Role of the Central Server

Our trust-based incentive mechanism relies on a central server. We assume that the central server is trustworthy. In this subsection, we discuss the important role of the central server, and believe that it is still valuable.

In a marketplace operating with our mechanism, the central server runs auctions to bring buyers and sellers together. It acts as a message relay station. In our setting, sellers can register to the central server with the information about the products they produce. The central server is similar to a registry in Paolucci et al. (2003) or a service broker mentioned in Maximilien and Singh (2004). For each request, it receives from a buyer, the central server forwards the request to the relevant sellers in the market. This avoids a lot of message overhead in the network. Sellers do not need to provide every buyer with the information about their products.

The central server also stores ratings of sellers provided by buyers. Using these ratings, it selects for each buyer a list of other buyers that are most trustworthy to the buyer as its neighbors. The trustworthiness of these other buyers is modeled by the central server for the buyer using the approach of Zhang and Cohen (2006) or the TRAVOS model proposed by Teacy et al. (2005). The central server also models the global reputation of each buyer based on the social network of buyers.

When the buyer models the trustworthiness of the sellers that register to join the buyer's auction, it may acquire from the central server ratings of the sellers provided by its neighbors. If a seller is allowed to bid to sell its product to the buyer, the central server also reports to the seller the reputation of buyers in the marketplace. The seller can then use this information to formulate its bid by giving more attractive offers to more reputable buyers.

After the transaction between the buyer and the selected winner of its auction is done, the buyer will report the result of conducting business with the seller to the central server, registering a rating for the seller. These ratings of the seller can then be shared with those buyers that consider this buyer as their neighbor, through the central server.

Each buyer can model the trustworthiness of other buyers without relying on the central server. However, we argue that it is still necessary to have the central server for the following reasons.

- The central server can assure that the ratings provided reflect a transaction that has actually occurred. In this way, there is some control over buyers trying to simply flood the system by untruthfully rating a seller a large number of times without any cost (referred to as "flooding" the system (Dellarocas 2000)).
- When a rating is fetched, these ratings are accurate because the central server always has up to date information. In a fully distributed approach, the ratings that are held by an agent may become stale.
- Third, if not relying on the central server, a buyer may untruthfully report the result of modeling the trustworthiness of other buyers. This will result in sellers' inaccurate estimation of buyers' reputation. In such a case, the overall mechanism is challenged. The honest buyers may not get appropriate rewards from sellers because their reputation is lowered by dishonest buyers. Similarly, dishonest buyers may get rewards. By relying on the central server, the trust that a buyer has of another buyer can be modeled by the central server. Thus, there is no need to ask for buyers' reporting of other buyers' trustworthiness.

6.4. Support from Studies in Economics

Our proposal of having sellers in the mechanism decrease their prices and increase quality of products to attract more buyers is also supported by theoretical work in economics. In Kranton (2003), Kranton argues that under the condition that a firm can permanently increase its market share by attracting new consumers with a price cut, it will have an incentive to produce high-quality goods. That is, its offer of high-quality goods will be credible, despite the lower current-period price. The profits from selling to a larger set of consumers in the future is greater than the one-shot gain from being dishonest and producing low-quality goods.

7. CONCLUSIONS AND FUTURE WORK

In this paper, we present a detailed incentive mechanism to encourage honesty, intended for use in designing e-marketplaces. We provide 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 validate 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.

Our particular proposal of allowing sellers to model buyers is promoting a novel direction within the artificial intelligence subfield of multiagent systems. With our proposal of designing an incentive mechanism based on trust modeling and also encouraging trust modeling researchers to consider incentives to diminish the problem of untruthful ratings, we hope to bridge research in the areas of trust modeling and mechanism design. Our ultimate goal is to contribute toward the design of effective electronic marketplaces populated by buying and selling agents. We aim to enable these agents to earn the trust of their users, as a result of our proposed methods for modeling the trustworthiness of their fellow agents and our proposed mechanism for promoting honest agent behavior. As such, our work should provide valuable encouragement for the acceptance of e-commerce by human users and business organizations.

In future work, we will investigate how a seller should learn an optimal value for the parameter ΔP_b , the increment in probability of being allowed to join a buyer's auctions in the future after the seller satisfies the buyer in a transaction. This is a part of the seller's estimation for expected future profit and the amount of rewards offered to the buyer (see equations (8) and (14)). The seller will lose profit if it sets ΔP_b too high because it will offer rewards that it cannot recover in sales. An optimal value for ΔP_b may be learned by the seller by examining its success in being accepted into auctions, to maximize the seller's profit over time.

We make several important assumptions in our current work to simplify our model. For example, we assume that feedback provided by buyers is objective, only indicating whether or not sellers delivered their promise in their bids. We also assume that buyers and sellers are rational. In future work, we will try to relax these assumptions to extend our model so that it can be applicable to many practical situations.

In our current work, a threshold k in Figure 3 is used for buyers to select the most trustworthy other buyers as their neighbors. Another threshold m in equation (5) is used

for buyers to allow only the most trustworthy sellers to join the buyers' auctions. In future work, we will investigate continuous functions for buyers to choose neighbors and for them to choose sellers to join their auctions. In addition, the work of Hazard and Singh (2010) provides some hints about how to discount buyer utility based on sellers' trustworthiness. We will investigate how their approach may provide the possibility of taking into account sellers' trustworthiness when selecting the winning seller in the auction.

We will also 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. Work by Kerr (2010) on coalition detection may be especially valuable to integrate. We will also analyze the scalability of our incentive mechanism in terms of, for example, the population of agents in marketplaces growing increasingly large or for a longer duration of operation, instead of 30 days, as in our current experiments.

We are also interested in operating with the "ideal" size of neighborhood. Toward this end, we are exploring how to limit the number of neighbors but to ensure that each one is able to offer sufficiently valuable advice, by using what is referred to as "advisor referrals," where each neighbor may in effect be replaced by one of its neighbors who is well respected and knowledgeable (Gorner and Cohen 2010). We will also analyze how the threshold k in Figure 3 used for limiting the number of neighbors affects the reputation of buyers and the effectiveness of our incentive mechanism.

REFERENCES

- BAR ISAAC, H. 2005. Imperfect competition and reputational commitment. *Economics Letters*, **89**(2):167–173.
- BRAYNOV, S., and T. SANDHOLM. 2002. Trust revelation in multiagent interaction. *In Proceedings of CHI'02 Workshop on the Philosophy and Design of Socially Adept Technologies*, Minneapolis, MN, pp. 57–60.
- CALZOLARI, G., and G. SPAGNOLO. 2006. Reputation and Collusion in Procurement. Manuscript, University of Bologna and Stockholm School of Economics.
- DASH, R. K., S. D. RAMCHURN, and N. R. JENNINGS. 2004. Trust-based mechanism design. *In Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multiagent Systems*, New York, pp. 748–755.
- DELLAROCAS, C. 2000. Immunizing online reputation reporting systems against unfair ratings and discriminatory behavior. *In Proceedings of the 2nd ACM Conference on Electronic Commerce (EC)*, Minneapolis, MN, pp. 150–157.
- DELLAROCAS, C. 2002. Goodwill hunting: An economically efficient online feedback mechanism for environments with variable product quality. In J. Padget et al., editors, *Agent-Mediated Electronic Commerce IV. Designing Mechanisms and Systems*, volume LNCS 2531. Springer-Verlag: Berlin, pp. 238–252.
- GINTIS, H., E. A. SMITH, and S. BOWLES. 2001. Costly signaling and cooperation. *Journal of Theoretical Biology*, **213**:103–119.
- GORNER, J., and R. COHEN. 2010. Optimizing advisor network size in a personalized trust-modelling framework for multi-agent systems. *In Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS) Workshop on Trust in Agent Societies (TRUST)*, Toronto, Canada.
- HAZARD, C. J. and M. P. SINGH. 2011. Intertemporal discount factors as a measure of trustworthiness in electronic commerce. *IEEE Transactions on Knowledge and Data Engineering*, **23**(5):699–712.
- JOHNSON, S., J. PRATT, and R. ZECKHAUSER. 1990. Efficiency despite mutually payoff-relevant private information: The finite case. *Econometrica*, **58**:873–900.
- JOSANG, A., and R. ISMAIL. 2002. The beta reputation system. *In Proceedings of the 15th Bled Electronic Commerce Conference*, Bled, Slovenia, pp. 324–337.

- JURCA, R., and B. FALTINGS. 2003. An incentive compatible reputation mechanism. *In Proceedings of the IEEE Conference on E-Commerce*, Newport Beach, CA, pp. 285–292.
- JURCA, R., and B. FALTINGS. 2004. CONFESS: An incentive compatible reputation mechanism for the online hotel booking industry. *In Proceedings of the IEEE Conference on E-Commerce*, San Diego, CA, pp. 205–212.
- JURCA, R., and B. FALTINGS. 2007. Collusion-resistant, incentive-compatible feedback payments. *In Proceedings of the 8th ACM EC*, San Diego, CA, pp. 200–209.
- KAMVAR, S. D., M. T. SCHLOSSER, and H. GARCIA-MOLINA. 2003. The eigentrust algorithm for reputation management in p2p networks. *In Proceedings of the 12th International Conference on World Wide Web (WWW)* Budapest, Hungary.
- KERR, R. 2010. Coalition detection and identification. *In Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, Toronto, Canada, pp. 1657–1658.
- KERR, R., and R. COHEN. 2007. Guaranteed security in trust and reputation systems. *In Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'07) Workshop on Trust in Agent Society*, Honolulu, HI, pp. 80–89.
- KIM, I.-G. 1998. A model of selective tendering: Does bidding competition deter opportunism by contractors? *The Quarterly Review of Economics and Finance*, **38**(4):907–925.
- KRANTON, R. E. 2003. Competition and the incentive to produce high quality. *Economica*, **70**(279):385–404.
- MAXIMILIEN, E. M., and M. P. SINGH. 2004. Toward autonomic web services trust and selection. *In Proceedings of the 2nd International Conference on Service Oriented Computing (ICSOC)*, New York, pp. 212–221.
- MILLER, N., P. RESNICK, and R. ZECKHAUSER. 2005. Eliciting informative feedback: The peer-prediction method. *Management Science*, **51**(9):1359–1373.
- PAOLUCCI, M., K. SYCARA, T. NISHIMURA, and N. SRINIVASAN. 2003. Toward a semantic web ecommerce. *In Proceedings of the 6th Conference on Business Information Systems (BIS)*, Colorado Springs, CO, pp. 153–161.
- PAPAIOANNOU, T. G., and G. D. STAMOULIS. 2005. An incentives' mechanism promoting truthful feedback in peer-to-peer systems. *In Proceedings of the IEEE/ACM CCGRID 2005 (Workshop on Global P2P Computing)*, Cardiff, UK, pp. 275–283.
- RICHARDSON, M., and P. DOMINGOS. 2002. Mining knowledge-sharing sites for viral marketing. *In Proceedings of KDD*, Edmonton, Canada, pp. 61–70.
- SHACHAT, J., and J. T. SWARTHOUT. 2006. Procurement auctions for differentiated goods. *Experimental Economics Center Working Paper Series*, Georgia State University.
- TEACY, W. T. L., J. PATEL, N. R. JENNINGS, and M. LUCK. 2005. Coping with inaccurate reputation sources: Experimental analysis of a probabilistic trust model. *In Proceedings of 4th International Autonomous Agents and Multiagent Systems*, pp. 997–1004.
- TRAN, T., and R. COHEN. 2004. Improving user satisfaction in agent-based electronic marketplaces by reputation modeling and adjustable product quality. *In Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, New York, pp. 828–835.
- WAN, Z., and D. R. BEIL. 2009. RFQ auctions with supplier qualification screening. *Journal of Operations Research*, **57**(4):934–949.
- WHITBY, A., A. JØSANG, and J. INDULSKA. 2005. Filtering out unfair ratings in bayesian reputation systems. *The Icfain Journal of Management Research*, **4**(2):48–64.
- YU, B., and M. P. SINGH. 2000. A social mechanism of reputation management in electronic communities. *In Proceedings of the 4th International Workshop on Cooperative Information Agents*, Boston, pp. 154–165.
- YU, B., and M. P. SINGH. 2003. Detecting deception in reputation management. *In Proceedings of International Autonomous Agents and Multi Agent Systems (AAMAS)*, pp. 73–80.
- ZHANG, J., and R. COHEN. 2006. Trusting advice from other buyers in e-marketplaces: The problem of unfair ratings. *In Proceedings of the 8th International Conference on Electronic Commerce (ICEC'06)*, Fredericton, Canada, pp. 225–234.