Theoretical Validation and Extended Experimental Support for a Trust-Based Incentive Mechanism for E-Marketplaces

Jie Zhang
School of Computer Science
University of Waterloo
Waterloo, ON, Canada
j44zhang@uwaterloo.ca

Kate Larson
School of Computer Science
University of Waterloo
Waterloo, ON, Canada
klarson@uwaterloo.ca

Robin Cohen
School of Computer Science
University of Waterloo
Waterloo, ON, Canada
rcohen@uwaterloo.ca

Abstract

In this paper, we revisit an incentive mechanism 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 develop precise formulations for the expected future profit sellers can realize when offering particular rewards for buyers and theoretically prove that with sellers calculating bids to sell to buyers according to our formulae: i) buyers will be better off honestly reporting seller ratings and ii) sellers are better off being honest, to earn better profit. We then revisit the proposed buyer strategy and theoretically show that limiting the sellers considered in each auction also promotes seller honesty. We ultimately focus on providing experimental evidence in defense of our incentive mechanism with the proposed seller and buyer strategies, including the added feature that buyers can come and go in the marketplace. As a result, we propose our framework as a valuable method for designing trustworthy e-marketplaces.

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.

In (Zhang & Cohen 2007a; 2007b), trust modeling is promoted as an important avenue for creating incentives for 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. (Gintis, Smith, & Bowles 2001) which argues that altruism in one context signals "quality" that is rewarded by increased opportunities in other contexts. Experimental evidence is 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 build on the model of (Zhang & Cohen 2007a; 2007b) and have buyers use a social network to model sellers and sellers model buyers to offer varying rewards. We analyze and refine the formulae used by sellers when reasoning about immediate and future profit. This leads to a precise proposal for bidding behavior and for rewards to buyers that we theoretically prove will cause both rational buyers and rational sellers to behave honestly.

We also revisit the buyer strategy to emphasize the value of limiting the number of sellers being considered at each auction, theoretically proving that limiting the number of sellers promotes seller honesty. Buyer behavior in the context of the proposed seller strategy is also illustrated through a detailed example.

We conclude with an extensive section of experimental results, using the revised seller and buyer strategies, operating in a marketplace where buyers may come and go and demonstrate the proposed framework as valuable for promoting honesty and for generating profit, both for buyers and for sellers.

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 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.

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.

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. Information about the neighborhoods to which the buyer belongs 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 their 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).

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.

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. 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 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.

Seller Strategy We discuss our mechanism in the context of the Request For Quote (RFQ) system (Shachat & Swarthout 2003). 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 $\{f_1, f_2, ..., f_n\}$, as well as a set of weights $\{w_1, w_2, ..., w_n\}$ 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 $\{s_1, s_2, ..., s_m\}$ ($m \geq 1$) allowed to join the auction are able to know the buyer's values of their products, which can be formalized as follows:

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

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. We define the seller's instant profit as follows:

$$U_{s_i} = P_{s_i} - C_{s_i} \tag{2}$$

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.

We now begin to refine the mechanism of (Zhang & Cohen 2007a; 2007b), 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 kinds of bids they should offer to buyers.

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

$$U_b = V_b - P_{s_i} \tag{3}$$

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

$$S_{s_i} = V_b - C_{s_i} \tag{4}$$

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

The seller whose surplus offer is the highest will win the auction. The RFQ auction then becomes a first-price sealed auction. As argued by Shachat and Swarthout (Shachat & Swarthout 2003), a symmetric Bayes-Nash equilibrium surplus offer function can be derived as follows:

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

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

where 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:

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

The seller in our mechanism also reasons about the expected future gain from winning the current auction. It takes into account the reputation of the buyer b. In our mechanism, each buyer in the marketplace has a fixed number of neighbors that the buyer trusts the most and from which it can ask advice about sellers. This forms a social network of buyers. A buyer that always provides truthful advice about sellers will be trusted by many other buyers and become the neighbor of them. This buyer will be considered reputable. Cooperating with reputable buyers will allow the seller to build its reputation and to be known as a trustworthy seller by many buyers in the marketplace. It will then be able to obtain more opportunities of doing business with buyers and to gain more profit in the future. Based on this analysis, we make the following claim:

Claim 1 *The expected future gain the seller s_i can earn increases with the number of other buyers considering the buyer b as one of their neighbors.*

We use R_b (reputation of b) to denote the number of other buyers considering b as their neighbor and $E_{s_i}(R_b)$ to denote the amount of the expected future gain. We then have the following inequality:

$$\frac{\partial [E_{s_i}(R_b)]}{\partial R_b} \geq 0 \tag{7}$$

Let us first consider a simplified scenario where sellers $\{s_1, s_2, ..., s_m\}$ have the same productivity. They have the same cost for producing the products that are valued equally by the buyer. In other words, we make the following assumption:

Assumption 1 *The distribution of S_{s_i}, $F()$ is a uniform distribution.*

Let us also assume that the seller's lowest realized surplus S_L for a transaction is 0. Equation 6 then can be simplified as follows:

$$\begin{aligned} P^*_{s_i} &= C_{s_i} + \frac{\int_{V_L - C_H}^{S_{s_i}} [F(x)]^{m-1} dx}{[F(S_{s_i})]^{m-1}} \\ &= C_{s_i} + \frac{\int_0^{S_{s_i}} (\frac{x}{S_H})^{m-1} dx}{(\frac{S_{s_i}}{S_H})^{m-1}} \\ &= C_{s_i} + \frac{\frac{x^m}{m(S_H)^{m-1}}\Big|_0^{S_{s_i}}}{(\frac{S_{s_i}}{S_H})^{m-1}} \\ &= C_{s_i} + \frac{\frac{(S_{s_i})^m}{m} - 0}{(S_{s_i})^{m-1}} \\ &= C_{s_i} + \frac{S_{s_i}}{m} \end{aligned} \tag{8}$$

From Equations 2, 3 and 4, we can see that the seller's realized surplus is in fact equal to the sum of the buyer and the seller's profit. Since the seller has expected future gain from winning the current auction, the seller's realized surplus S_{s_i} can then be changed as follows:

$$\begin{aligned} S'_{s_i} &= U_b + U_{s_i} + \lambda E_{s_i}(R_b) \\ &= V_b - C_{s_i} + \lambda E_{s_i}(R_b) \\ &= S_{s_i} + \lambda E_{s_i}(R_b) \end{aligned} \tag{9}$$

where $\lambda \in [0, 1]$ is a discounting factor.[2] 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:[3]

$$O^*_{s_i} = S_{s_i} + \lambda E_{s_i} - \frac{\int_{\lambda E_{s_i}}^{S'_{s_i}} [F(x)]^{m-1} dx}{[F(S'_{s_i})]^{m-1}} \tag{10}$$

From Equations 3, 4 and 10, we then can derive the modified equilibrium bidding function for the seller as follows:

$$\begin{aligned} P^*_{s_i} &= C_{s_i} - \lambda E_{s_i} + \frac{\int_{\lambda E_{s_i}}^{S'_{s_i}} [F(x)]^{m-1} dx}{[F(S'_{s_i})]^{m-1}} \\ &= C_{s_i} - \lambda E_{s_i} + \frac{\int_{\lambda E_{s_i}}^{S'_{s_i}} (\frac{x}{S_H})^{m-1} dx}{(\frac{S'_{s_i}}{S_H})^{m-1}} \\ &= C_{s_i} - \lambda E_{s_i} + \frac{\int_{\lambda E_{s_i}}^{S_{s_i}+\lambda E_{s_i}} (\frac{x}{S_H})^{m-1} dx}{(\frac{S_{s_i}+\lambda E_{s_i}}{S_H})^{m-1}} \\ &= C_{s_i} - \lambda E_{s_i} + \frac{\frac{x^m}{m(S_H)^{m-1}}\Big|_{\lambda E_{s_i}}^{S_{s_i}+\lambda E_{s_i}}}{(\frac{S_{s_i}+\lambda E_{s_i}}{S_H})^{m-1}} \\ &= C_{s_i} - \lambda E_{s_i} + \frac{\frac{(S_{s_i}+\lambda E_{s_i})^m}{m} - \frac{(\lambda E_{s_i})^m}{m}}{(S_{s_i} + \lambda E_{s_i})^{m-1}} \\ &= C_{s_i} - \lambda E_{s_i} + \frac{S_{s_i} + \lambda E_{s_i}}{m} - \frac{(\lambda E_{s_i})^m}{m(S_{s_i} + \lambda E_{s_i})^{m-1}} \\ &= C_{s_i} + \frac{S_{s_i}}{m} - \frac{1}{m}[\frac{(\lambda E_{s_i})^m}{(S_{s_i} + \lambda E_{s_i})^{m-1}} + (m-1)\lambda E_{s_i}] \end{aligned} \tag{11}$$

Comparing Equation 8 with Equation 11, we can see that the seller should offer the buyer reward $D_{s_i}(R_b)$ as follows:

$$D_{s_i}(R_b) = \frac{1}{m}[\frac{(\lambda E_{s_i})^m}{(S_{s_i} + \lambda E_{s_i})^{m-1}} + (m-1)\lambda E_{s_i}] \tag{12}$$

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 Claim 1, the seller's expected future gain $E_{s_i}(R_b)$ is a monotonically increasing function of R_b, the number of other buyers considering the buyer b as

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

[3]We replace $E_{s_i}(R_b)$ by E_{s_i} because of space limitations.

their neighbor. 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:

$$\frac{\partial D_{s_i}}{\partial R_b} = \frac{\partial\{\frac{1}{m}[\frac{(\lambda E_{s_i})^m}{(S_{s_i}+\lambda E_{s_i})^{m-1}} + (m-1)\lambda E_{s_i}]}{\partial R_b} \tag{13}$$

$$= \frac{1}{m}[\frac{\partial \frac{(\lambda E_{s_i})^m}{(S_{s_i}+\lambda E_{s_i})^{m-1}}}{\partial(\lambda E_{s_i})}\lambda\frac{\partial E_{s_i}}{\partial R_b} + (m-1)\lambda\frac{\partial E_{s_i}}{\partial R_b}]$$

$$= \frac{\lambda}{m}[\frac{\partial \frac{(\lambda E_{s_i})^m}{(S_{s_i}+\lambda E_{s_i})^{m-1}}}{\partial(\lambda E_{s_i})} + (m-1)]\frac{\partial E_{s_i}}{\partial R_b}$$

$$= \frac{\lambda}{m}[\frac{m(\lambda E_{s_i})^{m-1}(S_{s_i}+\lambda E_{s_i})^{m-1}}{(S_{s_i}+\lambda E_{s_i})^{2m-2}} + m - 1 - \frac{(m-1)(S_{s_i}+\lambda E_{s_i})^{m-2}(\lambda E_{s_i})^m}{(S_{s_i}+\lambda E_{s_i})^{2m-2}}]\frac{\partial E_{s_i}}{\partial R}$$

$$= \frac{\lambda}{m}[\frac{m(\lambda E_{s_i})^{m-1}}{(S_{s_i}+\lambda E_{s_i})^{m-1}} - \frac{(m-1)(\lambda E_{s_i})^m}{(S_{s_i}+\lambda E_{s_i})^m} + m - 1]\frac{\partial E_{s_i}}{\partial R_b}$$

$$\approx \{\underbrace{\frac{m(\lambda E_{s_i})^{m-1}}{(S_{s_i}+\lambda E_{s_i})^{m-1}}}_{\geq 0} + (m-1)\underbrace{[1-(\frac{\lambda E_{s_i}}{S_{s_i}+\lambda E_{s_i}})^m]}_{\geq 0}\}\frac{\partial E_{s_i}}{\partial R_b}$$

$$\geq 0$$

We have now proved the following proposition:

Proposition 1 *Under Assumption 1, sellers are better off providing better rewards to reputable buyers.*

The above analysis depends on Assumption 1. We can generalize this result to the case where sellers may not have the same productivity. 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:

$$P^*_{s_i} = V_b - S_{s_i} + \frac{\int_{S_L}^{S_{s_i}}[F(x)]^{m-1}dx}{[F(S_{s_i})]^{m-1}} \tag{14}$$

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

$$\frac{\partial P^*_{s_i}}{\partial S_{s_i}} = \frac{\partial\{V_b - S_{s_i} + \frac{\int_{S_L}^{S_{s_i}}[F(x)]^{m-1}dx}{[F(S_{s_i})]^{m-1}}\}}{\partial S_{s_i}} \tag{15}$$

$$= \frac{\frac{\partial[\int_0^{S_{s_i}}F(x)^{m-1}dx]}{\partial S_{s_i}}}{[F(S_{s_i})]^{m-1}} - \frac{\frac{\partial[F(S_{s_i})]^{m-1}}{\partial S_{s_i}}\int_0^{S_{s_i}}[F(x)]^{m-1}dx}{[F(S_{s_i})]^{2m-2}} - 1$$

$$= 1 - \frac{(m-1)\frac{\partial F(S_{s_i})}{\partial S_{s_i}}[F(S_{s_i})]^{m-2}\int_0^{S_{s_i}}[F(x)]^{m-1}dx}{[F(S_{s_i})]^{2m-2}} - 1$$

$$= -\frac{(m-1)\frac{\partial F(S_{s_i})}{\partial S_{s_i}}}{[F(S_{s_i})]^m}\int_0^{S_{s_i}}[F(x)]^{m-1}dx$$

$$\leq 0$$

Based on Equation 7, it is obvious that the seller's modified realized surplus S'_{s_i} formalized in Equation 9 will also increase with the increase of R_b:

$$\frac{\partial S'_{s_i}}{\partial R_b} = \frac{\partial[S_{s_i} + \lambda E_{s_i}(R_b)]}{\partial R_b} = \lambda\frac{\partial[E_{s_i}(R_b)]}{\partial R_b} \geq 0 \tag{16}$$

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 buyer that is considered as a neighbor by a greater number of other buyers 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 neighbors of many other buyers 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:

$$T_b = \sum_{k=1}^{l} U_{b,k} = \sum_{k=1}^{l}(V_{b,k} - P^*_{s_k}) \tag{17}$$

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:

$$\frac{\partial T_b}{\partial R_b} = \frac{\partial[\sum_{k=1}^{l}(V_{b,k} - P^*_{s_k})]}{\partial R_b} \tag{18}$$

$$= \sum_{k=1}^{l}\frac{\partial V_{b,k}}{\partial R_b} - \sum_{k=1}^{l}\frac{\partial P^*_{s_k}}{\partial R_b}$$

$$= -\sum_{k=1}^{l}\frac{\partial P^*_{s_k}}{\partial R_b}$$

$$\geq 0$$

since $\frac{\partial P^*_{s_k}}{\partial R_b}$ is negative (and considering $V_{b,k}$ as independent of R_b). Therefore, it is better off for the buyer to be honest and maintain high reputation, in order to gain more total profit.

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 the "Experimental Results" section. We also discuss how this strategy creates incentives for sellers to deliver what they promised in their 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 (Zhang & Cohen 2006) and the TRAVOS model proposed by Teacy et al. (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. A seller is considered trustworthy if its trust value is greater than a threshold γ. It will be considered untrustworthy if the trust value is less than δ. The buyer in our mechanism will allow only a number of the most trustworthy sellers to join the auction. If there are no trustworthy sellers, the sellers with trust values between γ and δ may also be allowed to join the auction.

However, buyers may provide untruthful ratings of sellers. Our mechanism allows the central server to maintain a fixed number[4] 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.[5] The trustworthiness of these other buyers (advisors) then also needs to be modeled. In the experiments presented in the "Experimental Results" section, the approach of Zhang and Cohen (Zhang & 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. 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. Finally, the trustworthiness of the advisor will be modeled by combining the private and public reputation values.

The approach (Zhang & Cohen 2006) for modeling the trustworthiness of sellers and advisors is "personalized", allowing buyers to set different weights for their personal experience and public knowledge. It also introduces the concept of a time window to discount older ratings and to avoid the situation where some advisors may provide a large number of untruthful ratings of sellers (known as "flooding").

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 states 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 (Calzolari & 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, in order to gain more profit in the future.

In (Kim 1998; Calzolari & 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. 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, as follows:

$$T_s = U_s = P_s - C_s \tag{19}$$

The studies of (Kim 1998; Calzolari & 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 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 E'_s is then greater than or equal to E_s, the seller's expected gain in their case. Greater expected future gain leads to a larger realized surplus (see Equation 9). Based on the argument supported by Equation 15 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:

$$T'_s = U'_s = P'_s - C_s \le P_s - C_s = T_s \tag{20}$$

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.*

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. In both examples, we also show that honesty sellers and honest buyers gain more total profit.

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 non-price features of the product p, the weight for each non-price feature and the information about the conversion from descriptive non-price feature values to numeric values are presented in Table 1.

[4]Note that exactly how to best choose this fixed number is left for future work.

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

Table 1: Buyers' Evaluation Criteria for p

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

The central server forwards b's request to the sellers in the marketplace. Five sellers $\{s_1, s_2, s_3, s_4, s_5\}$ are interested in selling their products to the buyer. The buyer first models the trustworthiness of the sellers. Only the sellers s_1, s_2 and s_3 that are trustworthy 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_1, s_2 and s_3 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 11 and their surplus offer for the buyer O_s^* calculated using Equation 10 are listed in Table 2. In this example, we assume that the sellers' expected future gain from winning the buyer's current auction is 1 and the discounting factor λ is 0.9.

Table 2: Sellers Bidding for b's Request

Seller	Cost	S_s	P_s^*	O_s^*
s_1	5	5	6.06	3.94
s_2	6	4	6.72	3.28
s_3	8	2	8.04	1.96

The buyer b will choose the seller that has the largest surplus offer as the winner of the auction. In this case, s_1 will be the winner. The buyer pays 6.06 to seller s_1. Later on, seller s_1 delivers the product. Suppose that the seller delivers the product with 3 year warranty in one 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_1 gains the instant profit, which can be calculated according to Equation 2 as follows:

$$U_{s_1} = P_{s_1} - C_{s_1} = 6.06 - 5 = 1.06$$

The sellers s_4 and s_5 that are not trustworthy do not gain any profit. Therefore, it is better off for sellers to be honest.

Seller Bidding for Buyers' Requests

In this example, we illustrate how a seller s models reputation of buyers and specifies its bids for buyers' requests according to their reputation values. Suppose that there are 6 buyers, $\{b_1, b_2, b_3, b_4, b_5, b_6\}$. They request the same product p with the same evaluation criteria presented in Table 1, which specifies the two non-price features of p, the weight for each non-price feature and the information about the conversion from descriptive values to numeric values.

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

Seller s needs to decide how to bid for each buyer's request. It models the reputation of each buyer. Assume that each buyer is allowed to have only 3 neighbors in this example. The neighbors of each buyer are listed in Table 3. We calculate each buyer's reputation represented by the number of its neighborhoods as follows:

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

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

According to the reputation of each buyer, seller s specifies its bid for each buyer's request. It produces different instantiations of the product p for different buyers. Table 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 4: Products Produced for Different Buyers

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

Table 5 lists the seller's amount of expected future gain $E_s(R_b)$ from selling the products to the buyers with different reputation values. We assume the discounting factor λ to be 1 and the number of bidders in each auction is 5. We also calculate the realized surplus S_s using Equation 9 and the seller's equilibrium bidding prices P_s^* according to Equation 11, as presented in Table 5.

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	0.1	0.3	0.4	0.5	0.5
S_s	2	2.1	2.3	2.4	2.5	2.5
P_s^*	1.4	1.33	3.18	3.11	8.04	8.04

According to Tables 4 and 5, we can calculate the profit gained by the buyers using Equation 3, as follows:

$$U_{b_1} = 1.6, \quad U_{b_2} = 1.67, \quad U_{b_3} = 1.82$$

$$U_{b_4} = 1.89, \quad U_{b_5} = 1.96, \quad U_{b_6} = 1.96$$

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

Experimental Results

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

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

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

Note that the main differences here compared to the experimental results of (Zhang & Cohen 2007b) are coping with buyers coming and going in the marketplace and employing the deeper reasoning strategies of the sellers and the buyers, outlined earlier in this paper.

Promoting Honesty

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

[6]All experimental results in the "Experimental Results" section are averaged over 500 rounds of the simulation.

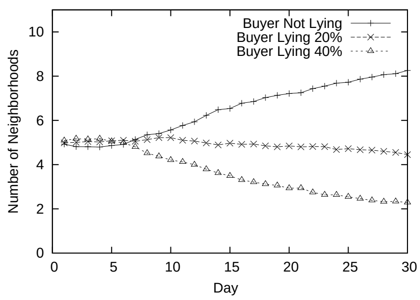

Figure 1: Reputation of Different Buyers

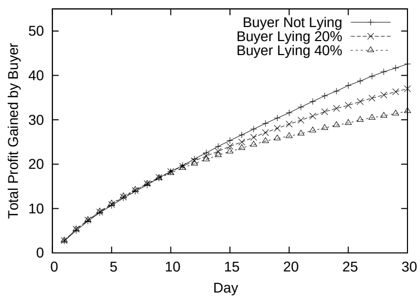

Figure 2: Profit Gained by Different Buyers

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 2, we can see that buyers providing fewer unfair ratings will gain more total profit. Note that the profit difference of different types of buyers is fairly 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.

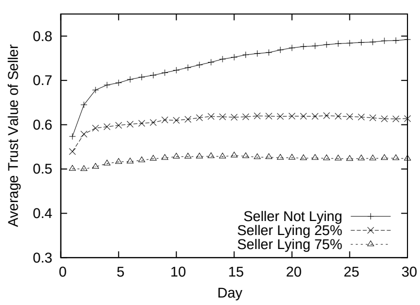

Figure 3: Average Trust Value of Different Sellers

We compare the average trust values of different sellers. The average trust value of a seller is calculated as the sum of the trust value each buyer has of the seller divided by the total number of buyers in the marketplace (90 in our experiments). As shown in Figure 3, results indicate that sellers being dishonest more often will have smaller average trust values. From this figure, we can see that the average trust values of the sellers being dishonest in 75% of their business are nearly 0.5.[7] 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.

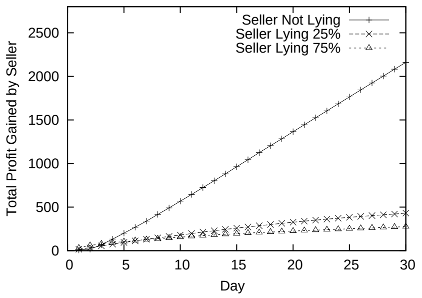

Figure 4: Total Profit Gained by Different Sellers

We also compare total profit gained by different sellers. Results are shown in Figure 4. From this figure, we can see that sellers being honest more often will gain more profit. Therefore, it is better off for sellers to be honest. We can also see that the profit difference between the honest sellers and the sellers lying 25% 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.

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.

We measure the average trust values of sellers from each group. Results shown in Figure 5 indicate that sellers modeling reputation of buyers will have higher average trust values. We also measure the total profit gained by different buyers. Results in Figure 6 indicate that sellers are better off to model reputation of buyers and adjust prices of products according to buyers' reputation, in order to gain more profit.

Buyer Strategy

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

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

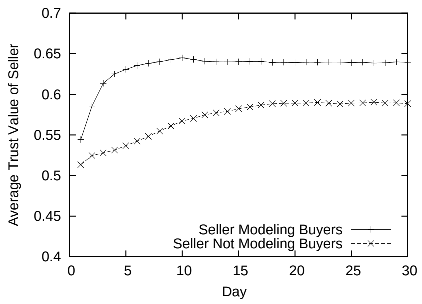

Figure 5: Average Trust Value of Different Sellers

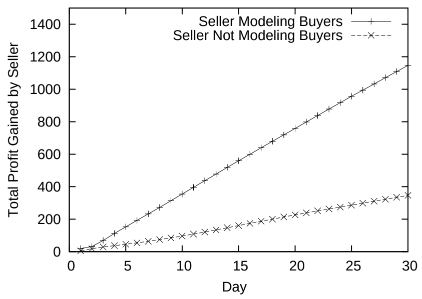

Figure 6: Total Profit Gained by Different Sellers

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

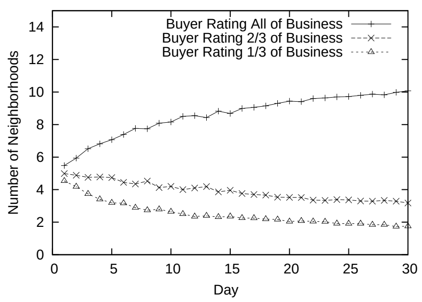

Figure 7: Reputation of Different Buyers

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

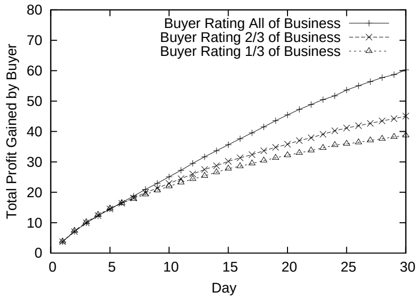

Figure 8: Profit Gained by Different Buyers

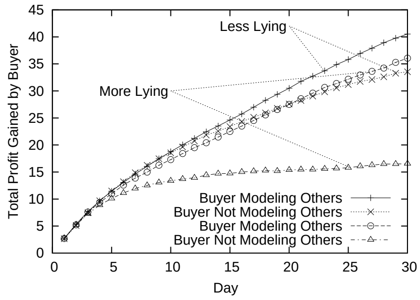

Figure 10: Profit Gained by Different Buyers

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

Buyer Modeling Sellers In this experiment, one third of the buyers models the trustworthiness of sellers based on their personal experience with the sellers 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.

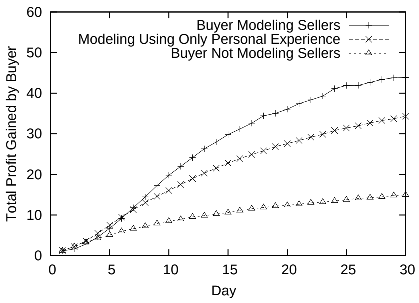

Figure 9: Profit Gained by Different Buyers

We compare the total profit gained by these three types of buyers. Results are shown in Figure 9. From this figure, we can see that buyers modeling the trustworthiness of sellers and 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.

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

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

Related Work

The approach adopted in this paper is to provide strategies for selling agents and buying agents in e-marketplaces to promote honest behavior by their business partners. This desired goal of creating marketplaces of trustworthy agents is also the aim of researchers focused on developing approaches for modeling trust and reputation. For example, the BRS system of (Whitby, Jøsang, & Indulska 2005) filters out ratings provided by advisors that are not in the majority amongst other ones, in a setting where probability density functions are used to estimate the reputation of a selling agent, propagating ratings provided by multiple advisors. The TRAVOS system of (Teacy *et al.* 2005) uses the approach of discounting the ratings provided by less trustworthy advisors.

While these methods can mitigate the effect of unreliable ratings, introducing direct incentives for honesty may be even more effective. For example, side payment mechanisms (Jurca & Faltings 2003; Miller, Resnick, & Zeckhauser 2005) 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.

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 (Jurca & Faltings 2003)). 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 (Jurca & Faltings 2007) 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 and detecting and avoiding dishonest, colluding sellers.

Conclusions and Future Work

In this paper, we presented a detailed incentive mechanism to encourage honesty, of use in designing e-marketplaces. For buyers, this is a matter of ensuring that they provide fair ratings of sellers. In our mechanism, a buyer maintains a neighbor list of trustworthy other buyers in order to model sellers and the sellers make use of this social network to model the reputation of buyers. Sellers then increase quality and/or decrease prices of products to buyers that are determined to be reputable. Buyers learn that they are better off providing truthful feedback when reporting ratings of sellers, thus becoming neighbors of as many other buyers as possible. Sellers are also kept honest, because buyers are modeling the sellers' trustworthiness, based on ratings provided by their trustworthy neighbors. We validated our mechanism through theoretical analysis and experiments.

In future work, we will look carefully into how sellers should model their expected future gain from winning the current auctions. A two-population evolutionary game theoretic model (Vytelingum, Cliff, & Jennings 2007) to analyze the complex interactions of buyers and sellers in marketplaces can be used here when estimating sellers' expected future gain. Another topic of future work is to determine the number of sellers allowed to join each buyer's auction, which ensures that dishonest sellers' instant profit does not exceed honest sellers' long-term profit. Kim (Kim 1998) provides some insights into how to derive an optimal number of bidders.

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, which has been well studied by Sen and Banerjee (Sen & Banerjee 2006). In addition, we are particularly interested in empirically demonstrating how our framework is able to handle marketplaces where strategic agents collude with each other.

References

Calzolari, G., and Spagnolo, G. 2006. *Reputation and Collusion in Procurement*. Manuscript, University of Bologna and Stockholm School of Economics.

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

Jurca, R., and Faltings, B. 2003. An incentive compatible reputation mechanism. In *Proceedings of the IEEE Conference on E-Commerce, Newport Beach, CA, USA*.

Jurca, R., and Faltings, B. 2007. Collusion-resistant, incentive-compatible feedback payments. In *Proceedings of the Eighth ACM Conference on Electronic Commerce (EC)*.

Kim, I.-G. 1998. A model of selective tendering: Does bidding competition deter opportunism by contractors? *The Quarterly Review of Economics and Finance* 38(4):907–925.

Miller, N.; Resnick, P.; and Zeckhauser, R. 2005. Eliciting informative feedback: The peer-prediction method. *Management Science* 51(9):1359–1373.

Sen, S., and Banerjee, D. 2006. Monopolizing markets by exploiting trust. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

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

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

Vytelingum, P.; Cliff, D.; and Jennings, N. R. 2007. Analysing buyers' and sellers' strategic interactions in marketplaces: An evolutionary game theoretic approach. In *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS) Workshop on Agent-Mediated Electronic Commerce*.

Whitby, A.; Jøsang, A.; and Indulska, J. 2005. Filtering out unfair ratings in bayesian reputation systems. *The Icfain Journal of Management Research* 48–64.

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

Zhang, J., and Cohen, R. 2007a. Design of a mechanism for promoting honesty in e-marketplaces. In *Proceedings of the Twenty-Second Conference on Artificial Intelligence (AAAI-07)*.

Zhang, J., and Cohen, R. 2007b. Towards more effective e-marketplaces: A novel incentive mechanism. In *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'07) Workshop on Trust in Agent Societies*.