Seller Bidding in a Trust-Based Incentive Mechanism for Dynamic E-Marketplaces

Jie Zhang and Robin Cohen

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

Abstract. In this paper, we develop a detailed bidding strategy for selling agents in electronic marketplaces, in a setting where buyers and sellers have incentives to be honest, due to a particular framework for trust modeling. In our mechanism, buyers model other buyers and select the most trustworthy ones as their neighbours 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 fair ratings for sellers, and are likely to be neighbours of many other buyers. Sellers will provide more attractive products to reputable buyers, in order to build their own reputation. We include simulations of a dynamic marketplace operating using our mechanism, where buyers and sellers may come and go, and show that greater profit can be realized both for buyers that are honest and sellers that are honest.

1 Introduction

People across the world today have embraced the Internet as part of their everyday life. While buyers and sellers can now find suitable business partners online, the promise of e-commerce will not be enjoyed unless electronic marketplaces are designed that provide users with some level of comfort that their partners can be trusted. Artificial intelligence provides that promise, by offering techniques from the traditional fields of user modeling and machine learning, in order for buyers and sellers to reason about each other.

In previous work [1], we promoted the use of a trust-based incentive mechanism to promote honesty in e-marketplaces populated by buying and selling agents. In particular, we observed that there are scenarios where buying agents in e-marketplaces would benefit from advice provided by other buying agents, when selecting the appropriate seller with which to do business. This may arise, for instance, when buyers may have limited experience with the population of sellers or in a scenario where buyers are migrating to different e-commerce environments in order to purchase goods, therefore failing to build up a longstanding history with sellers in any one environment. One major challenge, however, is the fact that these advisors may not always be truthful when providing ratings of sellers, offering unfairly high or unfairly low ratings (issues discussed in [2]).

S. Bergler (Ed.): Canadian AI 2008, LNAI 5032, pp. 368–379, 2008.

In the incentive mechanism that we propose, buyers are encouraged to be truthful in order to gain more profitable transactions. This idea is supported by Gintis et al. [3]. They argue that altruism in one context signals “quality” that is rewarded by increased opportunities in other contexts. In our mechanism, the reputation of buyers is modeled by sellers. A buyer is considered reputable if it is well respected in the community - i.e. it is a neighbour of many other buyers. This is also supported by Gintis et al. [3]. They argue that agents reporting honestly will be preferred by others as allies and will be able to attract a larger audience to witness their feedback. Sellers increase quality and decrease prices of products to satisfy reputable buyers, in order to build their own reputation. Our mechanism, therefore, creates incentives for buyers to provide fair ratings of sellers.

In this paper, we examine the seller strategy more clearly, specifying how sellers should bid, in order to make best use of our mechanism to enhance their reputability and therefore increase their profit. 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.

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

2 System Overview

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

In our system, a buyer that wants to purchase a product sends a request to the central server. This request indicates not only the product 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.

Each buyer maintains a neighbourhood of trusted other buyers, which will be asked to provide ratings of the sellers. As we will demonstrate in Section 4, it becomes very valuable to limit the number of sellers the buyer will consider for each auction, based on ratings it receives. 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 neighbour. Information about the neighbourhoods to which the buyer belongs is maintained by the central server and released to the sellers. Note that it is challenging for sellers to determine which bids to offer to buyers. We focus on this problem in the next section.

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 neighbour.

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 neighbours, choose a winning seller and report a final rating for the seller; sellers bid to win the sale to the buyer, consider buyer reputation in formulating their bids and then decide what product to deliver to the buyer (if at all).

3 Proposed Seller and Buyer Strategies

3.1 Seller Strategy

We discuss the seller strategy in the context of the Request For Quote (RFQ) system [4]. 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:

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

[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.

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.

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 [4], a symmetric Bayes-Nash equilibrium surplus offer function can be derived as follows:

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

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

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

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 neighbours that the buyer trusts the most and from which it can ask advice about sellers. This forms a social network of buyers. A buyer is reputable if it is the neighbour of many other buyers. Cooperating with reputable buyers will allow the seller to build its reputation and to be known as a trustworthy seller by many buyers in the marketplace. It will then be able to obtain more opportunities of doing business with buyers and to gain more profit in the future. We use R_b (reputation of b) to denote the number of other buyers considering b as their neighbor and $E_{s_i}(R_b)$ to denote the amount of the expected future gain. We then have the following inequality:

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

Let us consider a scenario where sellers $\{s_1, s_2, ..., s_m\}$ have the same productivity. They have the same cost for producing the products that are valued equally by the buyer. 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:

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

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:

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

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

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

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

We have already shown that sellers will gain better future profit when successful with more reputable buyers (Equation 7); this therefore suggests that the seller should offer better rewards to more reputable buyers as well. The bidding function outlined in Equation 11 provides for this, as the final term in the equation becomes a positive term times the change in E_{s_i}. Note that since the value of $P^*_{s_i}$ in Equation 11 is smaller than that of Equation 8 the reward for buyers can either be lower price or higher cost (i.e. greater product quality). The seller sacrifices its current profit in both cases.

3.2 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 [5] and the TRAVOS model proposed by Teacy et al. [6]. 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.

However, buyers may provide untruthful ratings of sellers. Our mechanism allows the central server to maintain a fixed number of neighbours 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.[3] The trustworthiness of these other buyers then also needs to be modeled. In the experiments presented in Section 4, the approach of Zhang and Cohen [5], combining personal experience and public knowledge is used for this purpose.

A final element of importance in the buyer's strategy is limiting the number of sellers being considered with each good that is being purchased. More specifically, the buyer will allow only a limited 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. Motivated by research from economics such as [7], this added restriction promotes honesty among sellers because honest sellers are offered sufficient future gain.

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

4 Experimental Results

We simulate a dynamic marketplace operating with buyer and seller strategies 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, that each buyer has a maximum of one request each day, and 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 neighbours.

There are also 9 sellers in total in the marketplace. Every 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. The sellers all follow our proposed bidding strategy.

4.1 Promoting Honesty

Here, we provide some general results to show that our proposed strategies promote 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 neighbour. The results[4] are shown in Figure 1(a). From this figure, we

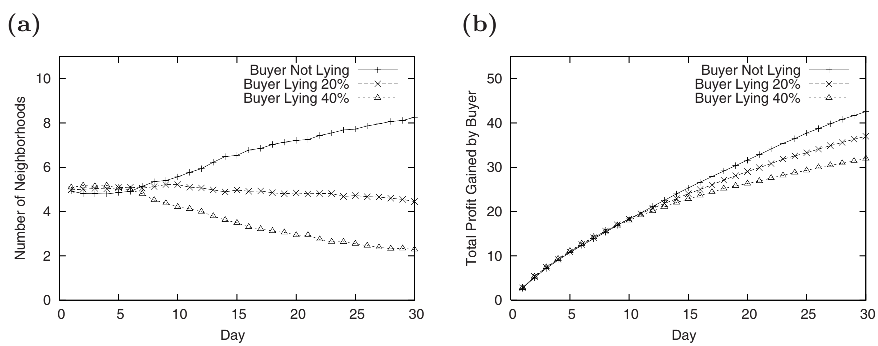

Fig. 1. Buyers' Reputation and Total Profit

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

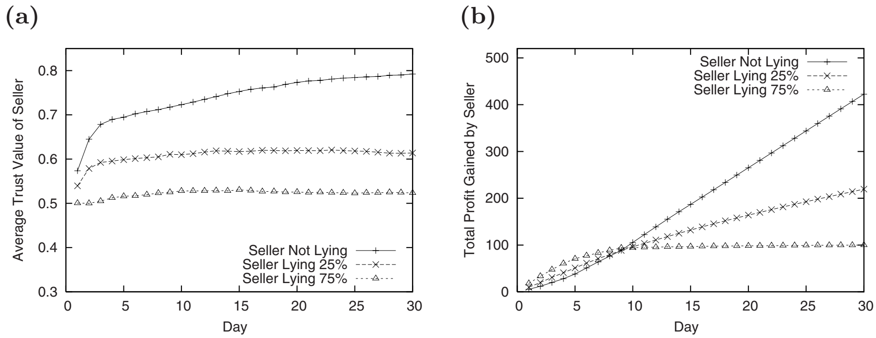

Fig. 2. Sellers' Average Trust and Total Profit

can see that the buyers providing the smaller percentages of unfair ratings will have the larger reputation values. Due to the randomness of the initial setting for our experiments, buyers' reputation values change stochastically at the beginning. After approximately 6 days when our marketplace converges, the changes of buyers' reputation will clearly follow a trend. After each day, we measure total profit gained by buyers that provide different percentages of unfair ratings. The profit gained by a buyer from buying a product is formalized in Equation 3. From Figure 1(b), 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.

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 2(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. 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 2(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 sellers lying more often may gain more profit in the first few days. When our marketplace converges, they will gain much less profit.

4.2 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

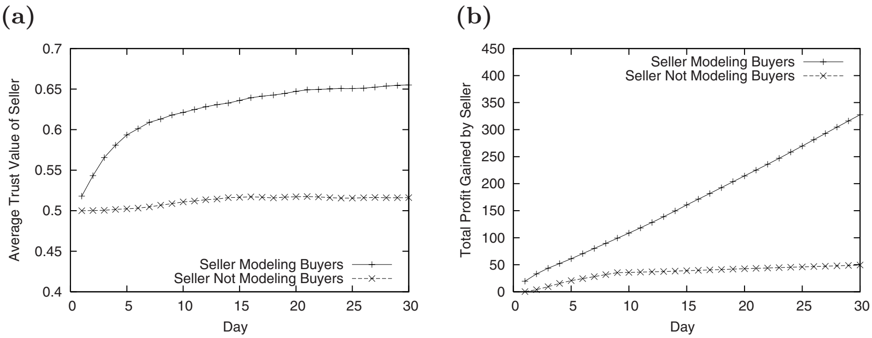

Fig. 3. Sellers' Average Trust and Total Profit

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 3(a) indicate that sellers modeling the reputation of buyers will have higher average trust values. We also measure the total profit gained by different buyers. Results in Figure 3(b) 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. Our proposed bidding strategy for sellers is shown to be effective.

4.3 Buyer Strategy

Limiting Number of Bidders. In the experiments in this section, we have 90 sellers. Similarly, every 30 sellers acts dishonestly in different percentages (0%, 25% and 75%) of their business with buyers. In the first experiment, we allow 30 sellers to join each buyer's auctions. Figure 4(a) shows the amount of

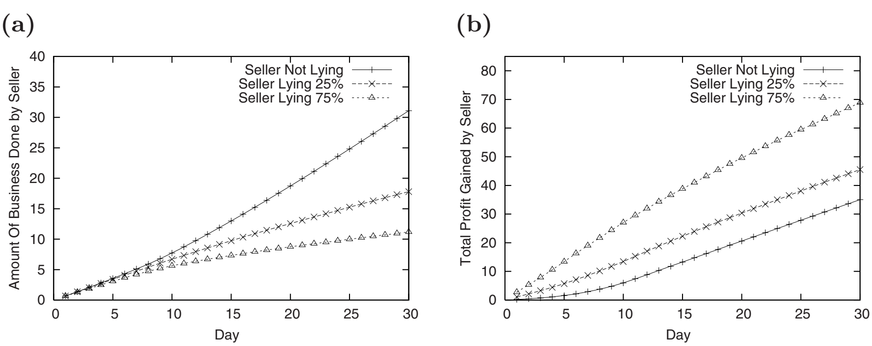

Fig. 4. Sellers' Amount of Business and Total Profit

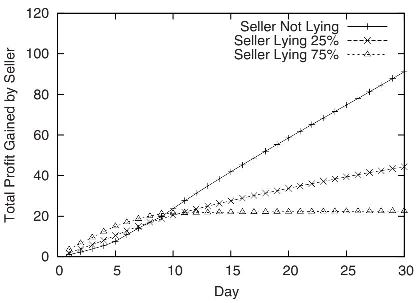

Fig. 5. Total Profit Gained by Sellers

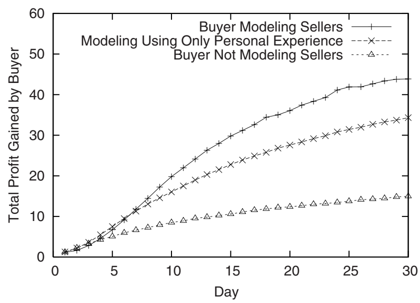

Fig. 6. Profit Gained by Different Buyers

business (number of transactions) done by different sellers. Sellers being honest more often are still able to gain more opportunities to do business with buyers. We also compare total profit gained by different sellers in this setting. However, from the results shown in Figure 4(b), we can see that sellers being dishonest more often will gain more total profit. In this case, because more sellers are allowed to join buyers' auctions, each seller's equilibrium bidding price should be lower in order to win the auctions. Sellers being honest gain very little profit from each business with a buyer; therefore, dishonesty will be promoted.

In the second experiment, we limit the number of bidders allowed in each of the buyers' auctions to be 6. As shown in Figure 5, sellers being honest more often are able to gain more total profit. Honest sellers in this case are more likely to win the future auctions of buyers. They are offered sufficient future gain because limiting the number of bidders increases each seller's equilibrium bidding price. Therefore, limiting the number of bidders will promote seller honesty.

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 neighbours. 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 randomly select some sellers to be allowed to submit bids.

We compare the total profit gained by these three types of buyers. Results are shown in Figure 6. 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 Related Work

The framework outlined in this paper has buyers modeling sellers, using ratings provided by advisors but also has sellers modeling buyers, in order to make effective bids in the marketplace. This is in contrast to the majority of approaches for modeling the trustworthiness of agents in e-marketplaces (e.g. the probabilistic reasoning model of TRAVOS [6]), that focus on methods for buyers to determine the reliability of advisors and hence of sellers. Tran and Cohen [8] do introduce seller modeling of buyers, but this framework focuses on direct experience alone and has sellers learning how to adjust quality and price of goods to satisfy buyer preferences. In contrast, our approach has sellers reasoning about how their reputation will be spread in the marketplace, leading to future gain, and models the reputability of the buyers.

We have also discussed the incentive for honesty among agents that results from our proposed buyer and seller strategies. A competing approach for creating incentives for honesty in e-marketplaces is the side-payment mechanism [9,10] that offers 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.

6 Conclusions and Future Work

In this paper, we proposed detailed bidding strategies for sellers and limits on sellers being considered by buyers, when using our trust-based incentive mechanism in e-marketplaces. Buyers acting as advisors learn that they are better off providing truthful feedback when reporting ratings of sellers, thus becoming neighbours 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 neighbours. With buyers limiting the number of sellers being considered when doing business, sellers are even more inclined to be honest, in order to maintain a profit. Our mechanism and our strategies are validated through experiments in a dynamic marketplace of significant size.

For future work, we will explore in greater detail how selling agents should formulate bidding strategies, when reasoning about competing agents in the marketplace. One promising approach is to estimate future gain using evolutionary game theory, as proposed in [11]. We should consider less uniform behaviour amongst the sellers as well. 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 [7] provides some insights into how to derive an optimal number of bidders.

We will also carry out more extensive experimentation in large-scale or real-world environments and continue to validate our model by comparing directly to models such as [9]. In our future experiments, we will also examine the situation where agents may vary their behaviour widely to exploit the marketplace, which has been well studied by Sen and Banerjee [12]. 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

1. Zhang, J., Cohen, R.: Design of a mechanism for promoting honesty in e-marketplaces. In: Proceedings of the Twenty-Second Conference on Artificial Intelligence (AAAI 2007) (2007)
2. Dellarocas, C.: Immunizing online reputation reporting systems against unfair ratings and discriminatory behavior. In: Proceedings of the 2nd ACM Conference on Electronic Commerce (EC), Minneapolis, MN (2000)
3. Gintis, H., Smith, E.A., Bowles, S.: Costly signaling and cooperation. Journal of Theoretical Biology 213, 103–119 (2001)
4. Shachat, J., Swarthout, J.T.: Procurement auctions for differentiated goods. Working Paper, IBM Research Labs (2003)
5. Zhang, J., Cohen, R.: A personalized approach to address unfair ratings in multiagent reputation systems. In: Proceedings of AAMAS 2006 Workshop on Trust in Agent Societies (2006)
6. Teacy, W.T.L., Patel, J., Jennings, N.R., Luck, M.: Coping with inaccurate reputation sources: Experimental analysis of a probabilistic trust model. In: Proceedings of AAMAS 2005 (2005)
7. Kim, I.G.: A model of selective tendering: Does bidding competition deter opportunism by contractors? The Quarterly Review of Economics and Finance 38(4), 907–925 (1998)
8. Tran, T., Cohen, R.: Improving user satisfaction in agent-based electronic marketplaces by reputation modeling and adjustable product quality. In: Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2004) (2004)
9. Jurca, R., Faltings, B.: An incentive compatible reputation mechanism. In: Proceedings of the IEEE Conference on E-Commerce, Newport Beach, CA, USA (2003)
10. Miller, N., Resnick, P., Zeckhauser, R.: Eliciting informative feedback: The peer-prediction method. Management Science 51(9), 1359–1373 (2005)
11. Vytelingum, P., Cliff, D., Jennings, N.R.: 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 (2007)
12. Sen, S., Banerjee, D.: Monopolizing markets by exploiting trust. In: Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS) (2006)