An Incentive Mechanism to Promote Honesty in E-marketplaces with Limited Inventory

Yuan Liu, Jie Zhang, and Qin Li

School of Computer Engineering
Nanyang Technological University, Singapore
{yliu3, zhangj, qin.li}@ntu.edu.sg

Abstract. In e-marketplaces with limited inventory where buyers' demand is larger than sellers' supply, promoting honesty raises new challenges: sellers may behave dishonestly because they can sell out all products without the necessity of gaining high reputation; buyers may provide untruthful ratings to mislead other buyers in order to have a higher chance to obtain the limited products. In this paper, we propose a novel incentive mechanism to promote buyer and seller honesty in such e-marketplaces. More specifically, our mechanism models both buyer and seller honesty. It offers higher prices to the products provided by honest sellers so that the sellers can gain larger utility. Honest buyers also have a higher chance to do business with honest sellers and are able to gain larger utility. Experimental results confirm that our mechanism promotes both buyer and seller honesty.

1 Introduction

In electronic marketplaces, lack of trust and reliability has been frequently cited to be one of the key factors that discourage buyers from participating. A reputation system, which predicts sellers' future behavior based on ratings given by buyers, is an effective way to help buyers to select good sellers [6]. It also creates incentives for sellers to behave honestly in order to be chosen by buyers. However, buyers may provide untruthful ratings to promote some sellers or drive some other sellers out of the market. To address this problem, incentive mechanisms, e.g. [5, 3], have been designed to supplement reputation systems, by creating an incentive for buyers to provide truthful ratings. One common but perhaps implicit assumption in these reputation systems and incentive mechanisms is that sellers can provide a large number of products in e-marketplaces. However, In the real world, e-marketplaces with limited inventory exist in many scenarios. One example is the hotel booking system for a famous tourism area during a peak season since booking a satisfactory hotel is often difficult. We call a marketplace in which the demand outweighs the supply *a marketplace with limited inventory*.

New challenges are imposed on promoting buyer and seller honesty in e-marketplaces with limited inventory. Sellers with limited inventory, given that other sellers also hold limited inventory compared to buyer demand, may behave maliciously in their transactions, by not delivering promised products or reducing the quality of delivered products. Even though their reputation would decrease due to the negative ratings from the buyers cheated by them, the sellers may still be willing to increase their profit by sacrificing reputation, because they may not have as a strong desire to maintain a very high reputation as in the marketplace where the supply outweighs the demand. Therefore, in the e-marketplaces with limited inventory, reputation itself cannot give sellers enough incentives to behave honestly. Buyers may also have incentives to report dishonest ratings. After a successful transaction with a seller, the buyer knows that the seller is a good seller. If the buyer provides a truthful (positive) rating about the seller, then the buyer reduces her own opportunity of doing business with the seller in the future, due to the limited inventory that the seller has. If the transaction is unsuccessful, reporting a truthful (negative) rating also reduces the buyer opportunity of doing business with other good sellers because other buyers will be less likely to do business with the bad seller but with the other good sellers, after taking the buyer's advice.

To address those challenges, we propose an incentive mechanism to promote buyer and seller honesty in e-marketplaces with limited inventory. In our mechanism, buyer honesty is measured by a *normalized proper scoring rule*, where a buyer can and only can gain maximal scores by providing truthful ratings. The higher score brings the buyer a higher expected utility. Seller honesty is measured by the ratings provided by buyers so that honest sellers are able to gain a high reputation. The products of sellers with a higher reputation are offered higher prices. This idea of the price premium is well supported by economic studies. Empirical evidence reveals that prices of products sold by honest sellers are generally higher [4]. The buyers' purchase intention would not be affected by the price premium provided to honest sellers [1]. Also, buyers with larger scores have more opportunities to conduct transactions with more reputable sellers. We conduct experiments to confirm that our mechanism promotes both buyer and seller honesty.

2 Our Incentive Mechanism

The e-marketplace employing our mechanism runs periodically. During each transaction period, each seller can only sell one product and each buyer can only buy one product. In the beginning of each transaction period, sellers post the products they want to sell and buyers post buying requests specifying the products they want to buy. The e-marketplace center gathers together the sellers who sell the same kind of products and the buyers who want to buy those products. It is assumed that in each transaction period, buyers' demand for the products is larger than sellers' supply of those products, meaning that the e-marketplace has limited inventory, and thus some buyers may not be able to do business with sellers. For the same products, their prices will then be determined by the e-marketplace center and these products will be allocated to some buyers. After each transaction, the buyer can provide a rating in $[0, 1]$ for the seller from whom the buyer receives the product, reflecting the buyer's satisfaction about the transaction, i.e. the ratio of the quality of the received product to that of the product promised by the seller. As the central component of the e-marketplace, our incentive mechanism is composed of a *normalized proper scoring rule*, a *reputation model*, a *pricing algorithm* and an *allocation algorithm*. More specifically, in our incentive mechanism, we measure buyer honesty by a *score* and seller honesty by *reputation*, which are updated after each transaction period. Buyer score will be updated after the buyer submits a rating according to the normalized proper scoring rule, making sure that truthful ratings provided by buyers could bring maximum scores. The seller reputation is calculated by the reputation model which aggregates ratings of the seller provided by buyers weighted by the scores of these buyers. The pricing algorithm sets higher prices for the products provided by sellers with higher reputation. The allocation algorithm ranks buyers according to their scores, and allocates products of honest sellers to buyers with the highest scores.

2.1 Modeling Buyer Honesty

Buyer honesty is measured as scores by *normalized proper scoring rules*. In this section, we provide a class of normalized proper scoring rules where buyers providing truthful ratings about sellers will be able to gain the maximal scores.

Given a binary event with two outcomes e and e', p is the actual probability of e and the actual probability of e' is $1 - p$. Let x be a predicted probability of e. If the outcome of the event is e, the agent having predicted the probability as x will be rewarded the scores $\mathbf{S}(x)$, while if the outcome is e', the agent will be rewarded $\mathbf{S}(1-x)$ scores. The expected amount of scores of the agent is denoted as $E(\mathbf{S}, x, p) = p\mathbf{S}(x) + (1 - p)\mathbf{S}(1 - x)$. The scoring function $\mathbf{S}(x)$ is a proper scoring rule, if and only if $E(\mathbf{S}, p, p) \geq E(\mathbf{S}, x, p)$ and the equality is true only when $x = p$ [2]. Based on the concept of proper scoring rules, we extend them to be normalized proper scoring rules, which are comparable, even when the scores are gained from the transactions with sellers having different honesty levels in delivering promised products.

Definition 1. *(Normalized Proper Scoring Rule $\mathbf{S}'$) Given a proper scoring rule $\mathbf{S}$, $Max(p) = \max_x E(\mathbf{S}, x, p)$ and $Min(p) = \min_x E(\mathbf{S}, x, p)$, a normalized proper scoring rule is defined as* $\mathbf{S}'(x) = \frac{\mathbf{S}(x) - Min(p)}{Max(p) - Min(p)}$.

From Definition 1, normalized proper scoring rules are bounded in $[0, 1]$. It is also essential that they have the same properties of the proper scoring rules, that is $E(\mathbf{S}', x, p) = p\mathbf{S}'(x)+(1-p)\mathbf{S}'(1-x)$, $E(\mathbf{S}', p, p) \geq E(\mathbf{S}', x, p)$, and equality is true only when $x = p$.

In our mechanism, the honesty of a seller s in delivering promised products is modeled by the seller's reputation R_s, which will be introduced in detail in the next section. Thus, the probability of s being dishonest is $1 - R_s$. In the end of the current transaction period t, a buyer b involved in the transaction with seller s can provide a rating indicating the buyer's satisfaction about the transaction. Once the rating is given, the buyer's score towards seller s measured by a normalized proper scoring rule as defined by Definition 1 will be updated. In consequence, the buyer's overall scores towards all sellers will also be updated.

Before we measure a buyer b's honesty $R_b(t)$, we first calculate the expectation value (denoted as $\overline{r}_b^s(t)$) of the distribution of the ratings provided by the buyer b towards seller s, including the rating given in the current transaction period. The buyer b's scores towards seller s can then be measured as follows:

$$R_b^s(t) = R_s(t-1)\mathbf{S}'(\overline{r}_b^s(t)) + (1-R_s(t-1))\mathbf{S}'(1-\overline{r}_b^s(t)) \qquad (1)$$

where $\mathbf{S}'$ is a normalized proper scoring rule and $R_s(t - 1)$ is the reputation of seller s up to the previous transaction. We also count the total number of ratings given by b

towards s, denoted as $N_b^s(t)$. By weighted averaging the scores gained towards different sellers, the buyer b's overall score is calculated as follows:

$$R_b(t) = \frac{\sum_{s\in S} R_b^s(t) \times N_b^s(t)}{\sum_{s\in S} N_b^s(t)} \tag{2}$$

where S is a set of all sellers whom the buyer b has done transactions with before and provided ratings for.

2.2 Modeling Seller Honesty

The honesty of a seller s is modeled by aggregating the ratings provided by buyers (who have previously conducted transactions with s) towards the seller s based on the respective buyers' scores reflecting the buyers' honesty in providing ratings. More formally, in the end of the transaction period t, given the expectation of the distribution of a buyer b's ratings $\overline{r}_b^s(t) \in [0, 1]$ towards seller s, buyer b's score $R_b(t)$ measured by Equation 2, and the number of transactions between buyer b and seller s denoted as $N_b^s(t)$, the reputation value (in $[0, 1]$) of seller s can be calculated as follows:

$$R_s(t) = \mathbf{F}(R_s(t-1), N_{b\in B}^s(t), R_{b\in B}(t-1), \overline{r}_{b\in B}^s(t)) \tag{3}$$

where B is a set of all buyers whom the seller s has done transactions with before and received ratings from, and $R_s(t-1)$ is seller reputation in the end of the previous transaction period $(t-1)$. **F** is a reputation model which can truly measure seller honesty in delivering promised product, and in this paper, we do not specify the form of **F**, since it is application dependent and many reputation modeling approaches have been proposed, such as [6].

2.3 Pricing and Allocating Products

In this section, we introduce the proposed pricing algorithm and allocation algorithm. For the purpose of simplicity, we focus on one kind of products[1], and assume that buyers' valuation of the products follows some distribution in the interval $[V_*, V^*]$ where V^* and V_* are the maximal and minimal valuation of buyers towards the products provided by sellers, respectively. We also assume that sellers have the same cost C of producing that same kind of products with the highest quality, and $V_* > C$, to make sure that honest sellers are profitable.

As we analyzed in the Section 1, sellers with limited inventory generally lack of the incentive to behave honestly even with reputation mechanisms employed because reputation information about sellers cannot impose competition among sellers in such markets, and sellers with relatively low reputation can still have the chance to do business with buyers because of the limited available products in the markets. The consequence is that sellers will decrease the quality of their delivered products (also reputation) to the point where buyers' utility is minimized (i.e. approaches 0) and at the same time maximize their own profit. In our mechanism, the pricing algorithm associates sellers' profit with their behavior. More specifically, it offers higher prices to products of sellers with

[1] Pricing and allocating is repeated for each kind of products.

higher reputation. In this way, it creates incentives for sellers to behave honestly. At the same time, the pricing algorithm makes sure that buyers can gain sufficient utility.

In our pricing algorithm, product prices are determined by a pricing function $\mathbf{P}(R)$, where R is seller reputation modeled by Equation 3. The pricing function should satisfy the some basic requirements: 1) $\mathbf{P}(R) > 0$ for $R \in (0, 1]$; 2) $\mathbf{P}(0) = 0$; 3) $\mathbf{P}(\delta) = C$; 4) $\frac{d\mathbf{P}(R)}{dR} > 0$; 5) $\mathbf{P}(R_0) = R_0 \times C$. Requirement 1 ensures that the price set for seller with positive reputation is larger than 0. In the extreme case where sellers never deliver products at all, the price for the sellers' products should be set 0 as in Requirement 2. In Requirement 3, δ is a reputation value set by our mechanism so that the price of products provided by sellers with reputation δ is exactly equal to C. Also, the price should increase with sellers' reputation (that is a monotonically increasing function), because sellers with higher reputation bear higher cost for delivering promised products. Since $\mathbf{P}(0) = 0$ and $\mathbf{P}(\delta) = C$, there should exist a reputation value R_0 so that $\mathbf{P}(R_0) = R_0C$, according to the continuity property of the pricing function $\mathbf{P}(R)$. Thus, when a seller's reputation $R = R_0$, the seller's profit would be $\mathbf{P}(R_0) - R_0C = 0$. In other words, R_0 is the minimum reputation with which sellers can gain non-negative profit. Sellers with reputation lower than R_0 will not be profitable. The purpose is to disappoint those sellers who intend to take advantages of the limited inventory situation by behaving dishonestly. By setting the lowest profitable reputation R_0, sellers with reputation lower than R_0 will generally leave the market.

To come up with a proper but simple pricing function, we started with a linear function for $\mathbf{P}(R)$, however it is impossible to satisfy all the basic requirements listed above. Thus, we choose a quadratic function in the general form $\mathbf{P}(R) = aR^2 + bR + c$. Given Requirement 2 ($\mathbf{P}(0) = 0$), we have $c = 0$. Given Requirements 3 and 5, we can derive $a = \frac{C(1-\delta)}{\delta(\delta-R_0)}$ and $b = \frac{C(\delta^2-R_0)}{\delta(\delta-R_0)}$. According to Requirement 4, we can also derive that $2aR + b > 0$, which can be satisfied by setting the constraint $\delta \geq \sqrt{R_0}$. The pseudo code summary of the pricing algorithm is shown in Algorithm 1.

Algorithm 1: The Pricing Algorithm

Input : S, a set of sellers offering the products;
R_s, reputation of a seller $s \in S$ before the current transaction period;
C, δ, R_0, which are introduced above;

Output : P, the price for a seller's product;

```
a = C(1−δ) / (δ(δ−R_0));
b = C(δ^2−R_0) / (δ(δ−R_0));
foreach s ∈ S do
    P_s = P(R_s) = a R_s^2 + b R_s;
```

In addition, our pricing algorithm has two nice properties. The first property is that buyers' profit is positive when R_0 and δ are set properly, ensuring that the buyers allocated with products of sellers will be willing to carry out the transactions with the sellers (see Proposition 2 in the next section). The second property is that buyers allocated products from sellers with higher reputation will be able to gain larger profit even though the prices of these products are higher. Therefore, buyers are willing to buy products from sellers with higher reputation (see Proposition 2). Due to the first property and the fact that not all buyers can be allocated with products (limited inventory), our allocation algorithm ensures that honest buyers (i.e. buyers with larger scores) will have higher probabilities of being allocated with products. Due to the second property, we make sure that honest buyers will also likely be allocated with products provided by sellers with higher reputation, so that honest buyers will be able to gain more profit. These create incentives for buyers to behave honestly by providing truthful ratings.

Algorithm 2: The Allocation Algorithm

Input : B, buyers who want to buy products;
S, a set of sellers offering the products;
η, the exploration factor;

Output : Allocation of products to some buyers;

1 $S_r \leftarrow$ Randomly choose η percentage of S (products);
2 $S_g \leftarrow$ The rest $1 - \eta$ percentage of S (products);
3 Sort S_g based on seller reputation in descending order;
4 Sort B based on buyer scores in descending order;
5 **foreach** $s \in S_g$ **do**
6   Allocate product of s to ranked top buyer $b \in B$;
7   Remove b from B;
8 **foreach** $s \in S_r$ **do**
9   Allocate product of s to random buyer $b \in B$;
10   Remove b from B;

Following the two properties of the pricing algorithm, we come up with the allocation algorithm whose pseudo code summary is shown in Algorithm 2. More specifically, the algorithm sets an exploration factor $\eta \in [0, 1]$. The η percentage of randomly selected products among all available products will be randomly allocated to some buyers (excluding the most honest buyers with the largest scores who will be allocated with another $1 - \eta$ percentage of products) (see Lines 8-10 in Algorithm 2). This is to make sure that new buyers will also have a fair chance to do business with sellers and later provide truthful ratings to gain scores. The η factor is set relatively high in the beginning of the operation of an e-marketplace when a large number of new buyers join the market, but will be decreased when the market becomes more mature and stable and not many new buyers will join the market. Another $1 - \eta$ percentage of all available products will be allocated to the most honest buyers (i.e. the buyers with the largest scores). in a greedy manner. To be specific, these products are sorted according to their sellers' reputation in a descending order. The buyers are also ranked in a descending order according to their scores. The products are then allocated to the buyers one by one according to the descending order, so that the products of sellers with higher reputation are given to the buyers with larger scores (see Lines 5-7 in Algorithm 2). Note that each buyer is allocated with one product in each transaction period.

3 Experimental Validation

In this section, we carry out a set of experiments to evaluate our incentive mechanism. We conduct our experiments in a dynamic setting. In the dynamic setting, some new sellers and buyers join the marketplace during the experiment.

We simulate an e-marketplace environment involving sellers and buyers exchanging the same kind of products. The total number of products provided by the sellers is less than that of the buyers' demand, i.e. a market with the limited inventory. We set $R_0 = 0.6$, $\delta = 0.85$ the cost in producing promised quality product $C = 1$, the minimal valuation of buyers towards the product $V_* = 2$ and maximal valuation of buyers towards the product $V^* = 2.5$, allocation exploration factor $\eta = 0.1$, reputation learning rate $\alpha = 0.5$, the maximal error rate of reputation model $\xi = 0.5$ and confidence level of reputation model $\gamma = 0.5$. Note that a set of simulations with variant settings has been experimented, and the results are similar.

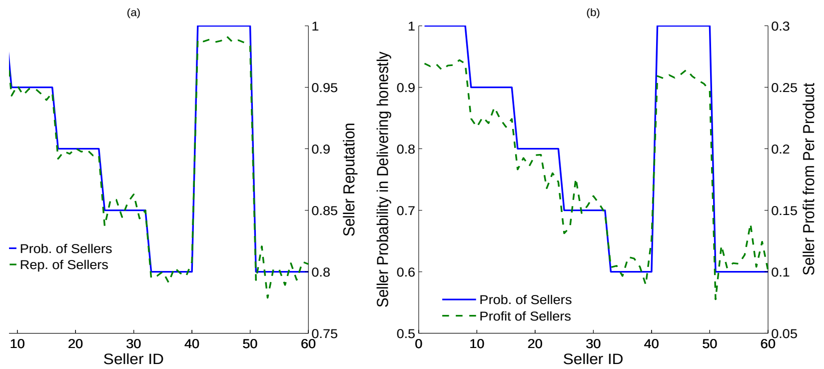

Fig. 1: The relationship between probability of sellers (a) seller reputation, (b) seller profit in selling one product. New buyers and sellers dynamically join the marketplace

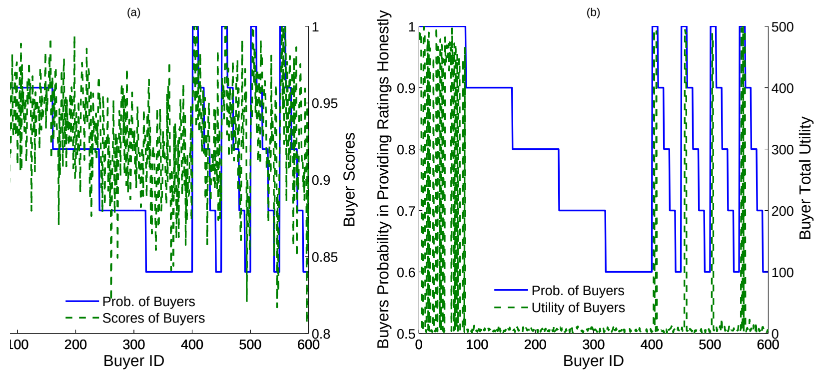

Fig. 2: The relation between buyer honesty and (a) buyer score, (b) buyer total utility

In our simulation, if a seller behaves honestly in one transaction, she delivers a quality product or a product with 50% quality. We set that the sellers have various probabilities in honest delivery and compare their average profit. For a buyer, if she behaves honestly, then she provides 1 for sellers who delivered quality products and 0.5 for sellers who deliver products with 50% quality. If the buyer is dishonest, then she provides 1 for sellers who delivered products with 50% quality and 0.5 for those who have delivered quality products. In the simulation, we allow new buyers and sellers join the marketplace during the simulation. In order to maintain our market constrain, i.e. e-marketplace with limited inventory, when a new seller joins, we allow 10 new buyers join into our system at the same time. After the boost-strapping, we let 5 new sellers and 50 new buyers (buyer honesty follows the same distribution with the existing 400 buyers) join into our simulation in every 100 transaction period. After 400 transaction periods, there are 20 new sellers (seller reputation follows the same distribution with the existing 80 sellers), and 200 new buyers participate into our market. After such a dynamic process, we simulate another 1000 transaction periods to observe seller profit and reputation. We obtain the results as shown in Figures 1 and 2.

In Figure 1, seller reputation and profit in selling one product (60 sellers in total) is shown. We observe that new honest sellers still gain the same reputation and profit as the sellers who previously existing in the e-marketplace. These results are shown in Figures 1(a) and 1(b), respectively. It means that honest sellers can always gain higher reputation and more profit no matter when they join our e-marketplace. In addition, more honest buyers gain higher scores and more utility, and these are shown in Figures 2(a) and 2(b). Therefore, the incentives of buyers and sellers in behaving honestly are still maintained when new sellers and buyers dynamically join into our e-marketplace. To conclude, our incentive mechanism ensures the sustainability of the e-marketplace by allowing new sellers and new buyers enter into our e-marketplace and our mechanism still works in such dynamic environment.

4 Conclusion

In this paper, we proposed an incentive mechanism to promote buyer and seller honesty in e-marketplaces with limited inventory. More specifically, a pricing algorithm is proposed to give high prices for products provided by honest sellers. In this way, sellers are incentivized to be honest. An allocation algorithm is proposed to allocate products of honest sellers to honest buyers. Conducting transactions with honest sellers will bring larger profit. Because of limited inventory, dishonest buyers may not be allocated any product. In this way, buyers are incentivized to be honest. We provide experimental verification for our mechanism.

References

1. Choe, Y.C., Park, J., Chung, M., Moon, J.: Effec of the food traceability system for building trust: Price premium and buying behavior. Information System Frontiers 11(2), 167–179 (2009)
2. Fang, F., Stinchcombe, M.B., Whinston, A.B.: Proper scoring rules with arbitrary value functions. Journal of Mathematical Economics 46(6), 1200–1210 (2010)
3. Jurca, R.: Truthful Reputation Mechanisms for Online Systems. Ph.D. thesis, EPFL (2007)
4. Mai, B., Menon, N.M., Sarkar, S.: No free lunch: Price premium for privacy seal-bearing vendors. Journal of Management Information Systems 27(2), 189–212 (2010)
5. Zhang, J., Cohen, R.: Design of a mechanism for promoting honesty in e-marketplaces. In: Procedings of the Conference on Artificial Intelligence (AAAI) (2007)
6. Zhang, J., Cohen, R.: Evaluating the Trustworthiness of Advice about Selling Agents in E-Marketplaces: A Personalized Approach. Electronic Commerce Research and Applications 7(3), 330–340 (2008)