An Incentive Mechanism for Eliciting Fair Ratings of Sellers in E-Marketplaces

Jie Zhang
David R. Cheriton School of Computer Science
University of Waterloo
Waterloo, ON, Canada N2L 3G1
j44zhang@uwaterloo.ca

1. RESEARCH PROGRAM

Our research is within the subfield of modeling trust and reputation in multi-agent systems for electronic commerce. More specifically, we are interested in addressing two problems that may arise in trust and reputation models where buying agents elicit opinions about selling agents from other buyers (known as advisors) in the marketplace:

- Unfair ratings of sellers provided to buyers
- Developing incentives for buyers to report their ratings of sellers

To explain, the ratings provided by advisors are possibly unfair. Buyers may provide unfairly high ratings to promote the seller. This is referred to as "ballot stuffing" [1]. Buyers may also provide unfairly low ratings, in order to cooperate with other sellers to drive a seller out of the marketplace. This is referred to as "bad-mouthing". Besides the problem of unfair ratings, rating submission is voluntary in most trust management systems. Buyers do not have direct incentives to provide ratings because, for example, providing reputation ratings of sellers requires some effort [3]. Providing fair ratings for a trustworthy seller may also decrease the chance of doing business with the seller because of competition from other buyers.

2. PROGRESS TO DATE

We first seek to develop a model that addresses unfair ratings. Our proposal is to adopt a personalized approach that allows a buyer to estimate the reputation (referred to as 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 the ratings for the sellers ever rated by the advisor. Finally, the trustworthiness of the advisor will be modeled by combining the weighted private and public reputations, where the weights are determined based on the estimated reliability of the private reputation, using probabilistic reasoning.

AAMAS'07 Doctoral Mentoring Program, May 13, 2007, Honolulu, Hawai'i, USA.

Equipped with the richer method for modeling trustworthiness of advisors in terms of private and public reputation, we are then interested in embedding this reasoning into a framework where there is as well incentive for being honest. Other researchers have also been working on developing incentive reputation mechanisms to encourage honesty in the reporting from buyers, in order to diminish concerns about unfair ratings. Two types of mechanisms have been developed, side payment mechanisms [4] and credibility mechanisms [5]. Side payment mechanisms offer side payment to buyers that fairly rate results of business with sellers. In these mechanisms, providing fair ratings for business results is a Nash equilibrium. Credibility mechanisms measure agents' credibility. The credibility of a buyer and a seller in a business will be decreased if their ratings about the business result are different.

We, however, begin with a novel insight that advisors may be motivated to provide honest ratings when asked by other buyers if advisors that are honest are rewarded by sellers through more profitable transactions. This idea is supported by the work in the field of evolutionary game theory, such as the work of Gintis et al. [2]. They argue that an agent's altruism in one context signals "quality" of the agent that will benefit from increased opportunities in other wider contexts. We use our personalized approach to create a social network of buyers. Each buyer in the society retains a neighborhood of the most trustworthy buyers, as advisors. In our mechanism, we also allow sellers to explicitly model the reputability of buyers, based on the number of neighborhoods they belong to in the society. Buyers that provide fair ratings of sellers are likely to be neighbors of many other buyers and can be considered reputable. This is also supported by Gintis et al. [2] through the model of a multi-player game. They argue that agents reporting honestly provide benefit to others and will further be preferred by others as allies. These agents will be able to attract a larger audience to witness their feedback (also known as increasing "broadcast efficiency"). Sellers in our system will increase quality and decrease prices of products for more reputable buyers, in order to build their own reputation. This therefore creates an incentive for buyers to provide fair ratings of sellers.

To date, we have developed a specific personalized model for representing the trustworthiness of advisors and sellers. One main idea that we use is to model the ratings that arrive according to their time windows. This helps to avoid the situation where advisors may untruthfully rate selling agents a large number of times and deal with changes of agents' behavior. Similarly, the personalized approach allows a buyer to model the private reputation of a seller based on its own ratings for the seller. If the buyer does not want to rely fully on its personal experience with the seller, it will ask for advisors' ratings of the seller. It then can derive a public reputation of the seller from ratings provided by them. The trustworthiness of the seller is modeled by combining the weighted private and public reputation values, using forgetting and discounting factors. We have carried out some experiments based on simulations to illustrate the effectiveness of our approach. For example, experimental results indicate that our approach can effectively model the trustworthiness of advisors even when buyers do not have much experience with sellers. Also, our approach is still effective when the majority of advisors provide large numbers of unfair ratings.

We have also begun the specification of the incentive mechanism. Consider the scenario in an electronic marketplace where a buyer B wants to buy a product p. We assume that the buying and selling process is operated as a procurement auction. The buyer B sends the request to a central server. The request contains information about a set of non-price features $\{f_1, f_2, ..., f_m\}$ of the product, as well as a set of weights $\{w_1, w_2, ..., w_m\}$ that correspond to how important each non-price feature is. The central server forwards the request to sellers in the marketplace. A seller $S \in \overline{S}$ sets the price and values for the non-price features of p. To gain profit from each possible transaction, the seller may not include in its bid the true cost of producing p with certain non-price features. The best potential gain the seller can offer the buyer from the transaction is as follows:

$$V'(p) = \sum_{i=1}^{m} w_i D(f_i) - C(p) \tag{1}$$

where $D()$ is a function to convert descriptive non-price feature values to numeric values and $C(p)$ is the cost for S to produce p. We define the distribution function for $V'(p)$ as $F(V')$. A symmetric Bayes-Nash equilibrium can be derived. The equilibrium bidding function can be derived as follows:

$$P^*(p) = C(p) + \frac{\int_{V_L - C_H}^{V'(p)} F(x)dx}{F(V')} - V_D(R) \tag{2}$$

where $V_D(R)$ is the valuation of discount for the buyer B with reputation $R(B)$, V_L is the lower bound of the value for the non-price features of p, and C_H $(V_L \geq C_H)$ is the higher bound of the cost for the seller to produce p.

Our mechanism allows the central server to maintain a fixed number of neighbors for each buyer from which the buyer can trust and ask advice about sellers. The central server models the trust value a buyer has of another buyer (an advisor) through the personalized approach. The seller S periodically acquires neighbor list information of buyers from the central server. It then counts for each buyer the number of neighborhoods. Suppose that there are N_B other buyers considering the buyer B as one of their neighbors. The reputation of B can be calculated as follows:

$$R(B) = \begin{cases} \frac{N_B}{\theta} & \text{if } N_B < \theta; \\ 1 & \text{otherwise.} \end{cases} \tag{3}$$

The value of θ depends on the total number of buyers in the marketplace. As can be seen from Equations 2 and 3, buyers that are neighbors of many other buyers will be offered more discount by sellers. Our mechanism also allows sellers to see how they have been rated by buyers, allowing sellers to reward those buyers deemed to be honest.

We have carried out preliminary experiments based on simulations to prove that both honest buyers and sellers are able to gain better profit in marketplaces operating with our mechanism.

3. FUTURE RESEARCH

Our research has two contributions, a personalized approach for buying agents to effectively model trustworthiness of other buyers and a novel incentive mechanism to elicit fair ratings of selling agents in electronic marketplaces. We are aware that many current social reputation models do not effectively allow for both public and private reputation modeling. For the future, we want to develop strategies for effectively comparing our model to other competing approaches. We may also learn more about how best to perform this modeling as we continue to make use of it for the problem of developing incentives for honesty in e-marketplaces.

For the incentive mechanism, one main direction for the future is to develop our mechanism in more detail. We will seek a more comprehensive approach for modeling buyers' reputation based on the social network topology. We are particularly interested in exploring how to demonstrate that our approach copes with collusion, whereas other incentive mechanisms do not (as noted by other researchers). Our mechanism allows sellers to view the ratings provided by buyers and can in this way detect dishonesty. It also allows buyers to maintain a list of trustworthy other buyers as their neighbors. If a buyer colludes, it can be excluded from neighborhoods and will not be rewarded by sellers. Sellers that collude will also not profit because buyers can make informed decisions about which sellers to do business with, based on advice from their neighbors. To prove the above expectations, we will develop experiments using agents that strategically collude. We seek to develop as well definitive comparisons with competing models. It may also be useful to determine how robust our model is to buyers and sellers leaving the marketplace.

4. REFERENCES

[1] C. Dellarocas. Immunizing online reputation reporting systems against unfair ratings and discriminatory behavior. In *Proceedings of the 2nd ACM Conference on Electronic Commerce*, Minneapolis, MN, October 2000.

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

[3] A. Jøsang, R. Ismail, and C. Boyd. A survey of trust and reputation systems for online service provision. *(to appear). Decision Support Systems*, 2005.

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

[5] R. Jurca and B. Faltings. Confess: An incentive compatible reputation mechanism for the online hotel booking industry. In *Proceedings of the IEEE Conference on E-Commerce, San Diego, CA, USA*, 2004.