

Sharing Models of Sellers amongst Buying Agents in Electronic Marketplaces

Kevin Regan¹, Thomas Tran², Robin Cohen¹

¹ School of Computer Science, University of Waterloo

² SITE, University of Ottawa

Abstract. In this paper, we demonstrate the value of decentralized models of sellers in electronic marketplaces, as the basis for purchasing decisions from buyers. We discuss how buying agents can model the reputation of sellers in order to make effective purchases and how these agents can also take advantage of reputation ratings provided by other buying agents in the marketplace, once it is established that each buyer will be independently modelling the sellers. We outline the methods required to make use of reputation ratings of sellers provided by other buyers, including adjustments for possibly different subjective scales and for possible deception in reporting the reputation ratings. In particular, we propose that buying agents model the reputation not only of sellers but of advisors, as part of the overall processing. We describe as well how buyers can most effectively model the set of possible buying advisors in order to apply selective processing of the information provided by these agents, in their purchasing decisions.

Keywords electronic marketplaces, reputation modeling, buying agents

Topic Areas modeling trust between users, agents modeling other agents, coalition formation based on affinity and trust

1 Introduction

In our research, we provide a framework for a buying agent in an electronic marketplace to make use of ratings of sellers provided by other buying agents, known as advisors. The solution is one where each individual agent maintains a separate model of the reputability of all selling agents and there is no merit in simply producing one centrally-held model of each seller. To the contrary, each agent will operate with different subjective standards in evaluating the sellers. Then when buying agents receive advice from advisors, they need to make adjustments for the differing subjectivity. In addition, in order to assist a buyer in determining which seller to choose for an upcoming purchase, it is in fact a subset of the different models of the seller that are brought together, on demand. While each buying agent may solicit advice from a set of possible advisors, the buyer will base its decision to purchase a good from a prospective seller solely on the basis of reputation ratings provided by reputable advisors. The remaining ratings that

are provided are used to adjust the models of the reputability of the advisors. As a result, we have a community of adaptive applications sharing information about possible sellers - clearly a scenario where decentralized modeling of agents is taking place.

In this paper, we first present a model for representing and adjusting the reputation ratings of sellers in electronic marketplaces. We then follow with a model for buyers to model the reputability of other buying agents within the marketplace, to possibly use them as advisors in purchasing decisions. In addition, we explain how a buyer can adjust both for the subjectivity of other buying agents and the possible deception generated by these buying agents, towards making purchases that will increase the satisfaction of the buyer. The key challenge that we will discuss in detail is: what methods can be used for a buyer to selectively choose a set of advisors from the entire set of other buyers in the marketplace. Part of the solution is to develop effective formulae for the adjustment of reputation ratings of advisors, subsequent to purchases being made, based on their advice. The overall aim of this research is to develop a model to drive purchasing decisions of buyers that can most effectively facilitate the sharing of information about potential sellers.

2 Model

2.1 Seller Reputation

In this section, we provide an overview of the model of Tran and Cohen [4]. The buyer primarily uses two criteria when selecting a seller. The buyer will use the reputation of a potential seller paired with an estimation of the value of the good that will be purchased.

Definition 1. *Given a set S of sellers. We denote the reputation of a seller $s \in S$ as seen by a buyer b as $r_s^b \in (-1, 1)$.*

For most of the properties discussed in this paper the use of superscript denotes who holds the property while the subscript denotes who it refers to.

Definition 2. *$f : G \times P \times S \rightarrow R$ is the estimated value function used by a buyer to assess the value of a good given the price and seller. We generally denote the estimated value function for a buyer b as $f^b(\cdot)$.*

For a seller that a buyer has no previous experience with, the reputation is initially set to zero and the estimated value is simply a function of the price of the good. It should be noted that because the initial reputation is significantly higher than the lowest possible reputation value, we have the potential problem of existing sellers with bad reputations re-entering the market using false pseudonyms. Some possible methods for eliminating this problem are discussed by Friedman and Resnick [3].

We use a reputation threshold Θ and a disreputation threshold θ to partition the set of sellers. Sellers for whom $r^b > \Theta$ are deemed reputable (R). Sellers for

whom $r^b < \theta$ are deemed disreputable (DR), while the rest of the sellers are put into the set (?)³ which the seller is unsure of. We can formally express this as follows

$$\forall s \in S \quad s \in \begin{cases} S_R^b & \text{if } r_s^b > \theta \\ S_{DR}^b & \text{if } r_s^b < \theta \\ S_?^b & \text{otherwise} \end{cases} \quad (1)$$

The reputation of a seller is adjusted based on resulting value of a transaction v^b and a buyer's satisfaction threshold ϑ^b . When $v^b \geq \vartheta^b$, the buyer is satisfied and the seller's reputation r_s^b is increased by $\mu(1 - r_s^b)$. When $v^b < \vartheta^b$, the buyer is unsatisfied and the seller's reputation is decreased by $\nu(1 - r_s^b)$. By setting $\nu > \mu$ we can ensure that reputation will be difficult to earn and easy to lose. We can also set μ and ν to be proportional to the amount of the transaction. Tran and Cohen [4] use

$$\mu = \frac{v^b - \vartheta^b}{\Delta v^b}, \quad \text{where } \Delta v^b = v_{max} - v_{min} \quad (2)$$

This method of adjustment provides two benefits. It allows the increase in reputation to be proportional to the value of the transaction. Thus, an automobile that was not delivered can impact a seller's reputation far more than a book that a buyer is unsatisfied with. The other benefit of this approach is that there is evidence that making reputation difficult to build and easy to tear down will discourage sellers from changing the value of their goods and allowing their reputation to oscillate between periods of building a reputation and periods in which they milk the reputation by over-advertising the quality of goods [2].

The buyer chooses the seller with the highest estimated value $f(\cdot)$ from among the reputable sellers. The potential sellers who have been deemed disreputable are never purchased from and the sellers a buyer is unsure of are occasionally used to buy goods from. The buyer selects a potential seller from the set $S_?^b \cup S_R^b$ with some small probability ρ in order to explore new sellers.

2.2 Advisor Reputation

We move beyond the model presented by Tran and Cohen [4] to provide an approach using seller ratings provided by other buyers.

Consider the situation after a buyer b has made a request for a good and received bids from a set S^p of potential sellers. In some situations it may be beneficial for the buyer to ask a set of other buyers about the potential sellers. For instance, when a buyer chooses a seller for the first time, or simply does not have much information about a seller it should consult other buyers. Also, if a buyer were to observe a radical change in behavior of a seller it had deemed reputable, then it could consult other buyers before returning to that seller. We refer to other buyers in this role as *advisors*. For each advisor $a \in A \subseteq B$ our

³ Tran and Cohen describe this set as those who are neither reputable nor disreputable

buyer will maintain a reputation r_a and partitions A_R , $A_?$, and A_{DR} in the same manner as seller information is maintained.

$$\forall a \in A \quad a \in \begin{cases} A_R^b & \text{if } r_a^b > \Theta' \\ A_{DR}^b & \text{if } r_a^b < \theta' \\ A_?^b & \text{otherwise} \end{cases} \quad (3)$$

The reputation of an advisor will be updated following a purchase when the buyer will either be satisfied or unsatisfied with the true quality of the good based on our satisfaction threshold ϑ . We essentially adjust the reputation of each as advisor based on whether they were right or wrong about the seller. Each seller reputation given by an advisor can be thought of as a prediction that categorizes a seller as either reputable or disreputable, with a third alternative being that the advisor is unsure. The direction of reputation adjustment for the advisor can be summarized with the following table

Result	Prediction	Change
Satisfied	Reputable	→ Increase
Satisfied	Disreputable	→ Decrease
Satisfied	Unsure	→ No change
Dissatisfied	Reputable	→ Decrease
Dissatisfied	Disreputable	→ Increase
Dissatisfied	Unsure	→ No change

We use the constant factors α and β to define the amount of the reputation adjustment. The adjusted reputation of an advisor a after an increase is defined as

$$\begin{aligned} r_a^b &\leftarrow r_a^b + \alpha(1 - r_a^b) \text{ if } r_a^b \geq 0, \text{ or} \\ r_a^b &\leftarrow r_a^b + \alpha(1 + r_a^b) \text{ if } r_a^b < 0 \end{aligned} \quad (4)$$

While the adjusted reputation of our advisor after a decrease is defined as

$$\begin{aligned} r_a^b &\leftarrow r_a^b + \beta(1 - r_a^b) \text{ if } r_a^b \geq 0, \text{ or} \\ r_a^b &\leftarrow r_a^b + \beta(1 + r_a^b) \text{ if } r_a^b < 0 \end{aligned} \quad (5)$$

In the preceding formulae α and β are positive and negative factors respectively and are chosen according to the preferences of each individual buyer.

After the adjustment of advisor's reputation, the advisors can be re-partitioned into reputable, unsure and disreputable sets using equation (3). This model of advisor reputation is used to decide which advisors to consult and how to interpret their feedback. For instance, a buying agent will avoid returning to advisors who have been moved into the disreputable set after an adjustment and only will only ask agents in the set of non-disreputable advisors (i.e. those in the set $A_R^b \cup A_?^b$) about a set S^a of sellers. The set S^a is composed of all the potential

sellers the buyer is unsure about as well as a set S_1^a of sellers which the buyer already knows about which is taken from $S_R^b \cup S_{DR}^b$. S_1^a will allow our buyer to assess how each advisor's standards differ and adjust in order to correct for them.

The advisor responses are combined to form a temporary reputation r_s^A for each seller. This new reputation is used to construct a set of reputable potential sellers (as in equation 1) from which the buyer can make a more informed purchase decision from among the reputable sellers. The way in which the advisor responses are combined must take into account the differing subjective standards used by each advisor to assess reputation as well as the possibility of the advisor being untruthful or inaccurate.

2.3 Advisor Subjectivity

To address the differing standards of an advisor, the buyer looks for any systematic difference in reputation and adjusts for it. As previously mentioned, our buyer asks each advisor about a set of sellers S_1^a that it already knows about through direct experience. We use this set of sellers to assess the similarity of the advisor to the buyer as follows

Definition 3. *For each advisor a that responds to the buyers request and seller $s \in S_1^a$, we may calculate the reputation error $\epsilon_s^a = r_s^a - r_s^b$*

If similar standards between agents were being used and the advisor was being honest, then the error ϵ_s^a would approach zero. However, if there is a systematic difference in the way an advisor determines reputation, then ϵ_s^a may be large, but would remain fairly constant over different sellers.

Definition 4. *We denote the mean and standard deviation of the reputation error over a set of sellers as $\bar{\epsilon}^a$ and σ^a respectively.*

To quantify this notion of how large ϵ_s^a is and how it varies, we find the mean $\bar{\epsilon}^a$ and standard deviation σ^a of ϵ_s^a across sellers. If σ^a were large, then ϵ_s^a does not represent any systematic difference and we cannot make an adjustment with any certainty. However, if σ^a is small, there is a systematic difference in the reputations that a has given b and we can adjust for this difference as follows

$$\forall s \in S^a, r_s^a \leftarrow r_s^a - \bar{\epsilon}^a \quad (6)$$

2.4 Advisor Deception

Our buyer will use the reputation held for each advisor to mitigate the effects of deceptive or inaccurate reputations given by an advisor. To avoid confusion between these two notions of reputation, we will occasionally refer to the reputation an advisor has about a seller as a prediction, since when this is information is passed on to the buyer and used as indirect reputation the advisors are, in a sense, making a prediction about the outcome of the buyer's purchase.

The responses from each of the advisors are combined so that the effect of dishonest sellers is minimized. However, each advisor is assumed to be honest until we find sufficient evidence of deception. It should be noted that we do not adopt the approach of weighing an advisor’s predictions by the advisor’s reputation ($r_a^b \cdot r_s^a$) that has been used by others [8, 6]. The argument for our approach is that until an advisor is no longer reputable, it is beneficial to fully consider their prediction (and not dilute it by some fractional weight).

We lessen the impact of dishonest sellers by maintaining reputations for each advisor and only use the predictions of the reputable advisors. We begin by finding the average over all the reputable advisors for each reputable seller.

Definition 5. *Given a seller s and a set of reputable advisors $A_R^b \subseteq A$, we denote the average prediction about s over all $a \in A_R^b$ as \bar{r}_s^A .*

An advisor with a high reputation who decides to lie about a particular seller can still have a large impact. This is particularly relevant since we assume all advisors are reputable until proven otherwise. To lessen the impact of reputable dishonest advisors we can choose to ignore predictions that are significantly different from that of the other reputable advisors. As a measure of significant difference we use the standard deviation of the prediction given by the reputable advisors, which we denote σ_s .

$$r_s^A \leftarrow \text{avg } r_s^a \text{ over } a \in A_R^b \text{ where } |r_s^a - \bar{r}_s^A| < \sigma_s \quad (7)$$

It should be noted that after a purchase a buyer’s reputation for *all* of the advisors contacted is updated. An advisor’s reputation can increase even if it was ignored when the seller was being chosen. In this way an advisor who fell below the reputable threshold can be redeemed.

3 Example

In this example we have only two potential sellers (s_r and s_{dr}) among whom our buyer b must decide to buy a good. The seller s_r has never deceived a customer, while s_{dr} has lied to customers. However, our buyer b , has no experience with either seller and seeks help from a set of advisors (a_1, a_2, a_3, a_4) with respective reputations (0.1, 0.5, 0.6, 0.4). Our buyer fixes the advisor reputation thresholds at $\Theta' = 0.20$ and $\theta' = -0.20$ resulting in a_1 being selected from the set A_r^b , while a_2, a_3 and a_4 are selected from the set A_{dr}^b . For the purposes of our example, a_1 turns out to be deceptive and provides deliberately inaccurate reputation information. The advisor a_2 is truthful, but has had good non-representative experiences with s_{dr} and provides an overly high reputation for this seller. Both a_2 and a_3 have high standards and this lowers the reputations they provide for each seller accordingly. a_4 is both truthful and has similar standards to our buyer.

Now, our buyer b receives a reputation for s_r, s_{dr} and $s_i \in S!$ from each advisor and if b were to simply average the reputations for s_r and s_{dr} without

Table 1. Example Details

	r_a^b	Explanation					
a_1	-0.1	deceptive, and not trusted					
a_2	0.4	truthful, high standards, inaccurate s_{dr} reputation	s_r	-0.25	-0.6	-0.7	0.2
a_3	0.5	truthful, high standards	s_{dr}	1.0	1.0	-1.0	-0.5
a_4	0.6	truthful, similar standards					

the methods developed to account for deception or differing standards, the result would be a reputation of -0.34 for s_r and 0.13 for s_{dr} . Now, let's say that b partitions sellers using: $\Theta = 0.20$ and $\theta = -0.20$ (as in equation 1), since $-0.33 < \theta$, s_r would be added to the set of disreputable sellers and since 0.13 is between θ and Θ , s_{dr} would be added to the set of sellers our buyer is unsure about.

The first step towards extracting accurate reputation information from our advisors is to account for any systematic bias. Our buyer finds the average difference between the reputation it holds and the reputation the advisor holds for each common seller $s'_i \in S_i$ ⁴. In the case of a_2 and a_3 , our buyer finds a difference of $\bar{\epsilon} = -1$ and a low σ indicating that our advisors consistently under-appreciate sellers by about -1. The buyer will adjust the reputations given by a_2 and a_3 by $-\bar{\epsilon}$. In our example 1 will be added to the reputations given by a_2 and a_3 and the average reputation for s_r and s_{dr} rises to 0.16 and 0.38 respectively⁵.

The second step is to ignore any reputation information from advisors that our buyer is unsure about. Here, the buyer ignores the deceptively low reputation that a_1 provided for s_r and the deceptively high reputation that a_1 provided for s_{dr} resulting in s_r 's reputation rising to 0.30 and s_{dr} 's reputation dropping to 0.17 . The seller s_r is now in our buyer's reputable set, however our buyer is still unsure about s_{dr} due to the inaccurate high reputation given by the truthful advisor a_2 .

The third and last step calculates the standard deviation of the set of reputations provided by reputable advisors and eliminates any reputation given by these reputable advisors that deviates from the average by more than one standard deviation. The unrepresentative high reputation provided for s_{dr} by a_2 is eliminated and the resulting average reputation for s_{dr} drops to -0.25 moving s_{dr} into the set of disreputable sellers. In our example the methods developed in this paper have successfully limited the effect of differing standards, and deceptive or inaccurate advisors. The effect of each phase of the combination process on the average reputation for each seller can be seen Figure 1. The buyer selects s_r and the reputation of the advisors is adjusted, depending on whether the buyer is satisfied with the purchase and the predictions of the advisors. After the interpretation phase our advisors a_1 , a_2 , a_3 , and a_4 have given a reputation of $(-0.25, 0.40, 0.30, 0.20)$ for s_r which predict s_r falling into the following respec-

⁴ The reputation ratings for each s'_i held by the buyer and advisor are omitted here

⁵ After adjustment a reputation greater than one will be normalized to one




Fig. 1. average reputation at each example step

tive sets (disreputable, reputable, reputable, reputable). Suppose that the buyer is satisfied with its purchase and our constant increase and decrease factors α and β are set to 0.2 and 0.4 respectively. This will result in the reputation of a_1 being decreased by $\beta \cdot (1 - r_{a_1}^b) = 0.36$, which will move a_1 into the set of disreputable advisors. The reputations of a_2 , a_3 and a_4 are increased to (0.52, 0.60, 0.68) respectively.

4 Related Work

Our approach to modeling reputation in electronic marketplaces can be contrasted with that of other researchers. Yu and Singh present a similar approach to trust in multi-agent systems [6, 7]. An agent builds a reputation for correspondents with which it interacts. If an agent has had no previous contact with a correspondent, it seeks out other agents to act as witnesses relating the reputation have established about that correspondent. As with our model, the reputation given by the witnesses is adjusted based on how well the witness was able to predict the reputation of other correspondents. They do not, however, take into account the fact that witnesses may be using different standards in determining the reputation of sellers.

In their most recent work [7] an agent takes the set of transactions with a correspondent and assigns each transaction to one of two sets based on two thresholds (much in the way our Θ and θ are used to categorize sellers). A transaction above the first threshold is considered evidence for trustworthiness, while a transaction below the other threshold is considered evidence against. The reputation of a correspondent is essentially a 3-tuple with the number of transactions giving evidence for, against and neither. While this approach successfully captures the uncertainty in reputation and how uncertainty gives way with new evidence, the evidence for or against trustworthiness is not weighted by the value of a transaction.

Our work also contrasts with that of Breban and Vassileva [1] where coalitions of agents are formed, based on trust. We model the reputability of all possible advisors and consult any of them who are not disreputable. In asking about sellers whose reputation is already known, we have a mechanism for determining the trustworthiness of these advisors. For future work, it would be useful to consider forming coalitions where advisors are motivated to remain trustworthy. This would integrate the important consideration of mechanism design [5] whereby careful determination of the “rules of interaction” can force a self interested buying agent to make decisions that respect the welfare of all buying agents.

Zacharia, Moukas and Maes have proposed a collaborative reputation mechanism for electronic marketplaces called Sporos which assigns each user a reputation and allows for the ratings of a group of users to be combined to form the reputation. Like our approach, the reputation of other users is taken into account when forming this reputation; however, Sporos weights each rating by the reputation of the user. Sporos addresses the problem of cheap pseudonyms by initializing the reputation of each new user to the lowest possible reputation. While this discourages existing users from re-entering the market as new users, it also unduly penalizes new users. The paper also presents a novel model for using chains of trust in highly connected networks, but does not address *when* other agents should be consulted, *what* criteria should be used to find these other agents, or *how* they should address the subjectivity in each agent’s ratings.

5 Discussion and Conclusions

This research presents a community of adaptive applications sharing user information, namely: a collection of buying agents, learning over time to make effective purchasing decisions, exchanging information about potential sellers. In fact, each buying agent models not only the sellers of the marketplace but also the other buyers, who serve as potential advisors. Our research presents strategies for determining how each agent should be representing its buyer and seller modeling information, so that when sharing is required, this sharing can be done in the most effective manner. In particular, buyers base their purchasing decisions on the reputation of sellers, and this reputation model is adjusted, due to information provided by buying advisors. But which buying advisors to consult and which advice to incorporate is carefully determined on the basis of models of those advisors’ deceptiveness and their subjectivity. Our decentralized seller model is in fact one that differs from buyer to buyer, according to the buyer’s demands and experiences. We think of each model in relation to its purpose (helping that buyer decide whether to buy from that seller).

As for how to ensure that different agents can share their information, we have a common representational framework: reputable and disreputable sets and thresholds for labelling sellers as belonging to these sets (specific to each individual buyer). We simply share reputation ratings between buyers and then make some adjustment for possibly differing standards between buyers. In order

to allow data harvested in one context to be useful for adaptation in another context, we propose receiving ratings from others and then, once more, adjusting for subjectivity. The “little player” can do this computation. As for how to represent the knowledge necessary to compute seller models on demand, this is simplified in our framework and falls out of the processing of each individual buying agent: it models the seller and adjusts reputation ratings and reputable sets after purchases.

One definite challenge is determining which advisors to ask for seller information. Our current proposal is to ask all buying advisors but to only “listen” to some of them. In addition, we consider asking those who have the best experience and who are most likely to be sharing similar standards. We provide a mechanism for interpreting data received by others in the marketplace, using formulae to adjust for subjectivity and deception. So, in our framework, the concept of trust arises precisely in this context.

As for applications, our model is designed to work in the context of electronic commerce and in particular in a marketplace driven by a contract net protocol set-up, where agents cannot determine whether they are content with goods until after purchases are made (and then adjustments to the user models need to be made). The topics we are most focused on are: modelling trust and similarity between users/agents and agents modeling other agents. We also have some contributions on the issue of coalition and community formation based on user affinity and trust.

References

1. Silvia Breban and Julita Vassileva. Using inter-agent trust relationships for efficient coalition formation. In *AI '02: Proceedings of the 15th Conference of the Canadian Society for Computational Studies of Intelligence on Advances in Artificial Intelligence*, pages 221–236, 2002.
2. C. Dellarocas. Immunizing online reputation reporting systems against unfair rating and discriminatory behavior. *E-Commerce*, October 2000.
3. E. Friedman and P. Resnick. The social cost of cheap pseudonyms. *Journal of Economics and Management Strategy*, 10:173–199, 1999.
4. T. Tran and R. Cohen. Improving user satisfaction in agent-based electronic marketplaces by reputation modelling and adjustable product quality. In *AAMAS04*, New York, USA, July 2004.
5. H. Varian. Economic mechanism design for computerized agents. In *Proceedings of the First USENIX Workshop on Electronic Commerce*, July 1995.
6. B. Yu and M. Singh. A social mechanism of reputation management in electronic communities. In *Proceedings of the 4th International Workshop on Cooperative Information Agents IV, The Future of Information Agents in Cyberspace*, pages 154–165, 2000.
7. B. Yu and M. Singh. Detecting deception in reputation management. In *AAMAS03*, pages 73–80, 2003.
8. G. Zacharia, A. Moukas, and P. Maes. Collaborative reputation mechanisms in electronic marketplaces. In *32nd Hawaii International Conference on System Sciences*, 1999.