Sharing Models of Sellers amongst Buying Agents in Electronic Marketplaces

Kevin Regan[1], Thomas Tran[2], Robin Cohen[1]

[1] School of Computer Science, University of Waterloo
[2] SITE, University of Ottawa

1 Introduction

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 all, we have a community of adaptive applications effectively sharing information about possible sellers.

2 Model

Our model builds on that that of Tran and Cohen [1], described briefly below.

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 $r_s^b \in (-1, 1)$.*

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

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

$$\forall\, s \in S \;\; s \in \left\{ S_R^b \text{ if } r_s^b > \Theta;\; S_{DR}^b \text{ if } r_s^b < \theta;\; S_?^b \text{ otherwise} \right\} \tag{1}$$

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

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

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 $\nu(1-r_s^b)$.

The buyer choses 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.

We move beyond the model presented by Tran and Cohen [1] 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 as seller it should consult other buyers. We refer to other buyers in this role as *advisors*. For each advisor $a \in A \subseteq B$ our buyer will maintain a reputation r_a and partitions A_R, $A_?$, and A_{DR} in the same manner as seller information is maintained.

$$\forall\, a \in A \quad a \in \begin{cases} A_R^b & \text{if } r_a^b > \Theta' \\ A_{DR}^b & \text{if } r_a^b < \theta' \\ A_?^b & \text{otherwise} \end{cases} \tag{2}$$

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 advisor based on whether they were right or wrong about the seller. There is an increase: if we were satisfied when the prediction was reputable, or if we were dissatisfied when the prediction was disreputable. There is a decrease if the satisfaction and reputability are at odds.

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

$$r_a^b \leftarrow r_a^b + \alpha(1 - r_a^b) \text{ if } r_a^b \geq 0;\ r_a^b \leftarrow r_a^b + \alpha(1 + r_a^b) \text{ if } r_a^b < 0 \tag{3}$$

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

$$r_a^b \leftarrow r_a^b + \beta(1 - r_a^b) \text{ if } r_a^b \geq 0;\ r_a^b \leftarrow r_a^b + \beta(1 + r_a^b) \text{ if } r_a^b < 0 \tag{4}$$

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 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 $S^a_!$ of sellers which the buyer already knows about which is taken from $S^b_R \cup S^b_{DR}$. $S^a_!$ 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^A_s 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.

Definition 3. *For each advisor a that responds to the buyers request and seller* $s \in S^a_!$*, we may calculate the reputation error* $\epsilon^a_s = r^a_s - r^b_s$

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.*

We can adjust for systematic differences (ie. σ^a is small) using the equation: $\forall\, s \in S^a,\;\; r^a_s \leftarrow r^a_s - \bar{\epsilon}^a$

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 ($r^b_a \cdot r^a_s$) that has been used by others [2]. The argument for our approach is that a 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* $A^b_R \subseteq A$*, we denote the average prediction about s over all* $a \in A^b_R$ *as* $\bar{r}^A_s$*.*

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. We then adjust seller reputation using the following rule: $r^A_s \leftarrow$ avg r^a_s over $a \in A^b_R$ where $|r^a_s - \bar{r}^A_s| < \sigma_s$.

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, 04). Our advisors a_1, a_2, a_3, a_4 provide respective seller reputations for s_r of $(-0.25, -0.6, -0.7, 0.2)$, and the respective seller reputations for s_{dr} of $(1.0, 1.0, -1.0, -0.5)$. Our buyer fixes the advisor reputation thresholds at $\Theta' = 0.20$ and $\theta' = -0.20$ resulting in a_1 being selected from the set $A^b_?$, while a_2, a_3 and a_4 are selected from the set A^b_R. 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 $s_! \in S_!$ from each advisor and if b were to simply average the reputations for s_r and s_{dr} without 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 $s'_i \in S_!$[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 respectively[5].

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

[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

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 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 respective 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 $\beta \cdot (1 - r^b_{a_1}) = 0.36$, 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 Discussion

This approach contrast with that of Yu and Singh [2], which appeals to advice from witnesses but does not account for differing standards in determining the reputation of sellers. It also compares to the Sporos system [3] which combines the ratings from a group of users, but does not consider how to find these other agents or how to address subjectivity.

Our research presents strategies for making decisions about sellers based on models of advisors' deceptiveness and subjectivity. This is a decentralized approach to representing sellers within the marketplace, where data harvested in one context is useful for adaptation in another, with each individual buyer managing its own processing. The modeling of trust between users and coalition formation based on trust are relevant issues within our framework.

References

1. 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.
2. B. Yu and M. Singh. Detecting deception in reputation management. In *AAMAS03*, pages 73–80, 2003.
3. G. Zacharia, A. Moukas, and P. Maes. Collaborative reputation mechanisms in electronic marketplaces. In *32nd Hawaii International Conference on System Sciences*, 1999.