Indirect Reputation Assessment in Electronic Markets

Kevin Regan *

Abstract

The Internet allows for A2A commerce at an unprecedented scale; anyone can do business with anyone. The new markets made possible by the Internet bring with them new challenges. This paper presents a system for buyers in electronic markets to avoid bad sellers by modeling the reputation of a seller. The model proposed by Cohen and Tran [6] is extended to provide a method for the exchange of indirect reputation information among buying agents. The subjectivity that arises when buyers use different standards to model seller reputation is addressed and a way to correct for any systematic differences between these reputations is developed. We assume that the indirect reputation shared by buyers may not be truthful and provide a model for the reputation of other buyers along with methods to minimize the impact of deceptive buyers. This work should be of interest to anyone wishes to address the issues of deception and cooperation in electronic markets.

1 Introduction

The Internet allows for A2A commerce at an unprecedented scale; anyone can do business with anyone. The new markets made possible by the Internet bring with them new challenges. How can we ensure that the users of our these new markets behave when they may have total anonymity and we lack the traditional real-world enforcement options to combat fraud such as fines and imprisonment?

Perhaps a natural approach would be to use cryptography and strong security protocols to limit what actions the users of a market can take. An example such an approach is using a Trusted Third Party as escrow services to ensure the delivery of a good before payment. However, added security comes at an added cost, in this case the fees charged by the escrow service and delays in receiving goods. Another approach is to use more flexible trust based systems to elicit good behavior in a market. The use of reputation for promoting trust has been given considerable attention in the context of centralized eBay-like markets [7, 8, 11], peer-to-peer networks [9, 12], and multi-agent systems [5, 6, 13, 14, 18, 19].

The centralized reputation systems used by Amazon.com [1] and eBay [2] aggregate buyer feedback to form a reputation for each seller. The advantage of this approach is that each buyer has access to a seller's global reputation, which can potentially take into account the experiences of every buyer that a seller has dealt with. However, since the marketplace is not directly involved in the delivery of goods, there is no simple way to verify the accuracy of buyer feedback.

At the other end of the spectrum are decentralized multi-agent systems which assess reputation and trust from the perspective of a single agent [6]. A buyer will model the reputation of each seller it has direct interactions with. Since the buyer is directly involved, there is no doubt about the outcome of transactions, however the buyer is limited to information about sellers it has already done business with.

This paper attempts to take advantage of the strengths of the multi-agent approach while providing an agent the possibility of learning about sellers they have not yet interacted with. This paper also draws inspiration from mechanism design which according to Varian [16]

> uses the tools of economics and game theory to design *rules of interaction* for economic transactions that will, in principle, yield some desired outcome.

Our desired outcome is modest in comparison with what may be the ideal outcome for a market in which every good is allocated to the agent that values it most and the welfare of all agents is maximized. Our goal is simply to provide the potential for this ideal market outcome by reducing behavior which is detrimental to the welfare of the market as a whole. We wish to prevent deceptive practices by sellers which may increase the utility of sellers but, by

*kmregan@uwaterloo.ca

lowering the utility of buyers, can damage the entire market. An example of this is a market in which we have some bad sellers who over-advertise the value of their goods (or simply do not deliver goods at all). If buyers cannot distinguish between good and bad sellers, many will leave the market for one in which they can expect more utility. As a result of buyers leaving the market the good sellers may no longer find it profitable and leave too. Eventually what we are left with is a market of lemons [4] consisting of only bad sellers.

We present a mechanism system for buyers to share their subjective reputations with the goal of avoiding dishonest sellers and buyers. We use the terms *direct* reputation and *indirect* reputation to distinguish between the seller reputation a buyer gains from direct interaction and the seller reputation a buyer receives from another buyer.

We will first get a better sense of the market the agents will be using through the use of some examples. We will then list the important challenges of avoiding bad sellers in these markets using a multi-agent approach. The paper will then present the single-agent model given by Cohen and Tran [6] and proceeds to extend this model by providing a system by which agents can ask other agents for their assessment of a seller's reputation. This will be followed up by a discussion of the benefits of this approach and a comparison to other approaches.

2 Some examples

The system presented by this paper is embedded in an electronic market with agents in fixed roles as either buyers or sellers. A buyer is free to purchase a good from any seller, however a seller may deceive a buyer by not delivering a promised good or by over-advertising the quality of the good. Purchases are negotiated with a three step process similar to that of Contract Net [15]. A buyer issues a request for particular good, the sellers who can provide this good make a bid at a certain price, and the buyer chooses between these bids and selects a seller.

An example of such a market would be that of computer parts. Imagine a set of small businesses who buy parts to assemble computers from a wide range of computer parts suppliers. The suppliers, or sellers, have an abundance of certain kinds of parts. One seller may specialize in displays while another in memory for laptops. Some newly manufactured computer parts will simply fail to work when installed after purchase and the varying quality of the sellers goods will reflect this. It should be noted that some groups of buyers will most likely end up doing business with similar sellers. For instance the makers of laptop computers will tend to purchase from sellers providing specialized laptop components. It should also be mentioned that the standards of the small businesses will vary and some low-cost computer sellers may not care as much about quality as other computer sellers with higher prices and extended warranties and service.

Another example which highlights the subjective aspect of reputation is that of a market for clothing. We would have a set of consumers buying a wide variety of clothing from many sellers. A buyer may want a pair of jeans which will vary in price and quality depending on which seller the buyer chooses. There is most likely an abundance of any one type of clothing, so consumers are not competing for that one t-shirt. The value each buyer places in an item of clothing is a matter of individual preference and as a result the reputations each buyer establishes for a seller will be highly subjective. However, one can imagine groups of buyers who share the same tastes and frequent the same sellers. One group may be interested in designer clothing and another in low-cost business attire.

3 Challenges

Any multi-agent system for electronic marketplaces with the aim of reducing transactions with bad sellers by having buyers communicate must overcome the following issues:

Trust Agents are self-interested and may be deceptive. Sellers may over-advertise the quality of goods or simply not deliver the goods after payment. When exchanging information about sellers, buyers may lie.

Subjectivity Buyers may have different standards upon which they base their reputation for sellers. Buyers receiving seller reputation information from other buyers must take this into account. A buyer accustomed to buying designer jeans may be unsatisfied with a pair of discount jeans. Another buyer who is generally satisfied with discount clothing could interpret a reputation given by the first buyer and adjust for their bias.

Cheap Pseudonyms We would like to make it easy for new sellers to enter the marketplace, however any seller whose reputation is worse than that of a new seller could simply enter the marketplace under a new identity. It is natural to assign a reputation to extremely dishonest sellers which is worse than that of sellers who are new, but to do this we need to make it difficult to discard a bad reputation and start anew with a false identity.

Cooperation We need an incentive for it to be in the best interest of each individual buyer to cooperate with each other and share information about sellers.

Efficiency Our solution should also take into account communication efficiency whenever possible to rule out things like each buyer communicating with every other buyer.

4 Model

We use the model presented by Cohen and Tran [6] as a starting point because it provides a way to partition sellers into a set the agent has deemed reputable, a set the agent has deemed disreputable and a set of sellers about which the agent is still unsure. The model can function as originally intended when the agent knows about the other sellers and can be naturally extended to include indirect reputation when an agent is unsure. We begin with a closer look at the original single agent model.

4.1 Direct Reputation

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 *$f : G \times P \times S \to \mathbb{R}$ 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 and the good (not taking the seller into account). 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 will be explained in the discussion section of this paper.

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 (?)[1] which the seller is unsure of. We can formally express this as follows

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

4.1.1 Adjustment of reputation

The reputation of a seller is adjusted based on resulting value of a transaction v^b and a buyers satisfaction threshold[2] 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. For instance we could use

$$\mu = \frac{v^b - \vartheta^b}{\Delta v^b}, \quad \text{where } \Delta v^b \text{ is the range of values} \tag{2}$$

[1]Cohen and Tran use the term non-reputable to refer to this set
[2]Denoted demanded product value by Cohen and Tran

This method of adjustment provides two benefits. It allows the increase in reputation to be proportional to the value of the transaction. Thus, an expensive suit that was not delivered can impact a sellers reputation far more than a pair of socks a consumer 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 it by over-advertising the quality of goods [8].

4.1.2 Selecting a seller

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 seller from the set $S_? \cap S^p$ with some small probability p in order to explore new sellers.

4.2 Indirect Reputation

We now move beyond the model presented by Cohen and Tran [6] to provide an approach to the use of indirect seller reputation 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. We refer to other buyers in this role as *advisors.* For each advisor $a \in B$ our buyer will maintain a reputation r_a and partitions B_R, $B_?$, and B_{DR} in the same manner as seller information is maintained.

Definition 3 *We denote the set of advisors used by a buyer b to evaluate a set of potential sellers S^p as $A_{S^p}^b$*

The model of our buyer needs to address the following questions: *when* should other buyers be queried, *who* should be queried, *what* should they be asked, and *how* the responses should be used. The next section will provide an overview of the way in which the model addresses each of these questions.

4.2.1 When

There is communication and processing overhead involved when asking advice of other buyers. To reduce this overhead the buyer will not seek the help of advisors when there is a set of reputable potential sellers to choose from. However, in many cases the set of potential sellers will not include any sellers that have been judged reputable. For instance, in our computer component example a buyer may decide to build a different kind of computer using new components which are not offered by familiar suppliers. Formally, the buyer will decide to ask for advice in the case that $S_R \cap S^p = \phi$.

Table 1: Advisor's Response

	s_1	$\dots$	s_n	s'_1	$\dots$	s'_n
b_1						
b_2						
$\vdots$		$\ddots$			$\ddots$	
b_m						

4.2.2 Who

Before giving the details of how we select the group of advisors, we first list some possible desiderata that we would like our advisors to satisfy.

Experience We would like to chose advisors who have had many interactions with the a large portion of the potential sellers.

Familiarity We would like to use an advisor repeatedly to build trust in the reputations provided by the advisor. If we initially choose our advisors in a naive way, they may not have experience with future potential sellers and we create a tradeoff between familiarity and experience.

Similarity We would like our advisors to have created reputations using standards that are similar to ours (for instance using a similar satisfaction threshold ϑ).

Proximity Others have suggested that selecting advisors based on how close they are in the network topology [12]. While this method will improve communication efficiency, it does not provide any guarantee that the advisors selected will have any experience with potential sellers, or similarity to the buyer.

Adopting an approach for selecting an advisor based on some of the above desideratum may preclude us from satisfying others. However, choosing advisors using similarity and familiarity will often lead to advisors who have experience with the potential sellers our buyer is interested in. Take for example our clothing market and a consumer who is purchases mainly designer clothing. By selecting a set of advisors who have similar standards and purchase designer clothing, it seems reasonable to assume that the new sellers that the advisors decide to buy from may overlap with the with future potential sellers that our buyer may wish to evaluate.

The details of how similar advisors are initially found will be fully explained in the section 5.4 after we have formalized the notion of similarity and show how it can be used to interpret the subjective reputations provided by the advisors. Essentially when deciding which advisors to ask, the buyer will consult its *advisor cache* which will store the set of advisors whom the buyer has used before as well as: the sellers that advisors has experience with, the reputation of the advisor and some measures of similarity between the advisor and the buyer. If a sufficient number of advisors having experience with potential sellers is not found in the advisor cache, new advisors could be added using one of the following approaches.

A request for new advisors can be propagated among buyers in the manner of a query in peer-to-peer networks like Gnutella [3]. When a buyer receives a request for advice about a set of potential sellers that have some overlap with its experience it would reply with a list of those sellers. Alternately, we could centralize this process using a *matchmaker* server. This sever would maintain an entry for each buyer along with a list of the sellers they have experience with. The buyers would periodically update these entries and occasionally send requests containing a set of potential sellers. These requests would be answered with a set of potential advisors and the sellers each advisor has experience with. These are only two possible approaches and a full exploration of these and other suitable approaches will be left to future work.

4.2.3 What

We ask our set A of buying advisors about the a set S^A of sellers that the buyer wishes to know about. 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 $S^b_R \cup S^b_{DR}$. The advisor cache contains a list of the sellers each advisor has experience with and can be used to construct $S^A_!$ (in a manner specific to each advisor).

In response the buyer will receive a set of reputations for each seller from each buyer that can be represented in using a matrix where $b_1 \dots b_m \in A$, $S^A_? = s_1 \dots s_n \in S_?$ and $S^A_! = s'_1 \dots s'_n \in S_R \cup S_{DR}$.

4.2.4 How

The specifics of how the user responses are used will be fully explained in section 6, but we can begin here with an overview. The way in which the buyer uses the responses from the advisors will take into account any subjective bias held by the advisor as well as the reputation the buyer holds for the advisor. This process can be decomposed into 3 distinct phases as illustrated in figure 1.

First the buyer *interprets* the responses of each advisor to correct for any bias, then the buyer choses the set of advisors to listen to and for each seller the reputations are *combined*. The result of the combination step is an advisor reputation r^A_s for each seller $s \in S^p$. The advisor reputation is used to again partition the set of potential sellers

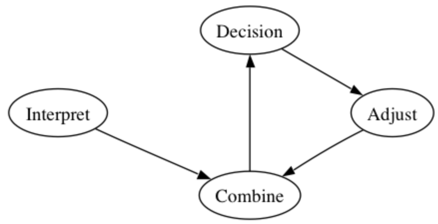

Figure 1: Cycle of Use

into those who are reputable, disreputable and those we still unsure of. The buyer can now use the procedure from Cohen and Tran's original model (as described in section 4.1.2) to choose the seller and make a purchase. Once the purchase has been made, the buyer uses the outcome to *adjust* the reputations of the advisors.

5 Some specifics

We now present a more detailed look at how the three phases of interpretation, combination and adjustment affect how the advisors responses are used.

5.1 Interpretation

The interpretation phase addresses the subjectivity inherent in the reputation information given by advisors. Our 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 $S_!^A$ 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 4 *For each advisor $a \in A$ and seller $s \in S_!^A$ we may calculate the reputation error $\epsilon_s^a = r_s^a - r_s^b$*

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

Definition 5 *We denote the mean of the reputation error over a set of sellers as $\bar{\epsilon}^a$.*

Definition 6 *We denote the standard deviation of the reputation error over a set of sellers as σ^a.*

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^a across sellers. 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

$$\forall\, s \in S^A,\;\; r_s^a \leftarrow r_s^a - \bar{\epsilon}^a \tag{3}$$

To illustrate how this interpretation phase might work we can return to our example of clothing market and envision a scenario in which two consumers, Allan and Bob, are interested mainly in business attire and purchase from an overlapping set of sellers. Bob asks Allan to advice him about a set of sellers S^A and the reputations Allan responds with fairly low, because Allan is very picky about his clothing and the smallest flaw will leave him unsatisfied. Bob, on the other hand may notice such flaws, but it does not taint the reputation he holds for the sellers as much as it does for Allan. Upon comparing the reputation of the sellers they have in common Bob notices that Allan does indeed have lower reputations for most sellers, but the reputations tend to be lower by the same error $\bar{\epsilon}^a$. Bob can now adjust the reputation by this error to compensate for Allan's high standards.

5.2 Combination

In this phase we make use of the reputation the buyer has about advisors as well as the reputation those advisors have about a seller. 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 in a sense the advisors are making a prediction about the outcome of the buyer's purchase.

The responses from each of the advisors is 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 advisors reputation ($r_a^b \cdot r_s^a$) that has been used by the Sporas system [19] and others [12, 17]. 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 7 *Given a seller s and a set of reputable advisors $A \subseteq A^b$, we denote the the average over all $a \in A$ 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 a reputable

Table 2: Combination example

advisor	a	c	d	e	f
reputation	0.6	0.1	-1.0	0.8	0.75

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 reputations of a seller given by the reputable advisors, which we denote σ_s.

$$r_s^A \leftarrow \text{ avg } r_s^a \text{ over } a \in A_R \text{ where } |r_s^a - \bar{r}_s^A| < \sigma_s \quad (4)$$

It should be noted that in the adjustment phase the buyers reputation for *all* of the advisors 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.

To illustrate the combination phase we walk through a scene involving our market for computer parts. There is a buyer b who asks the advisors a, c, d, e, and f about a set of potential sellers. They report the following reputations for a seller s with a real reputation of 0.75

All of the advisors are reputable except for c who has been wrong often enough that his reputation has fallen below the reputation threshold. Our buyer b calculates $\bar{r}_s^A \approx 0.28$ and finds that the reputation given by the deceptive advisor d is beyond a standard deviation[3] from $\bar{r}_s^A$ and ignores this prediction. The average $r_s^A = 0.716$ is then calculated using a, e, and f. If the c and d had been used the average would have been 0.25 which would most likely have been below the buyers reputable threshold of Θ.

5.3 Adjustment

Once we have calculated r_s^A for each seller we can partition them and choose the seller with the highest estimated value from among the reputable sellers. After the purchase has been made the buyer will either be satisfied or unsatisfied with the true quality of the good based on our satisfaction threshold ϑ and the reputation of the seller will be adjusted. We also adjust the reputation of each as advisor, essentially, based on whether they were right or wrong. The following table enumerates the change in the advisors reputation, based on the result of the transaction and whether the seller reputation given by the advisor would categorize the seller as reputable, disreputable or neither from the perspective of our buyer.

Satisfied	Reputable	$\rightarrow$	Increase
Satisfied	Disreputable	$\rightarrow$	Decrease
Satisfied	Unsure	$\rightarrow$	No change
Dissatisfied	Reputable	$\rightarrow$	Decrease
Dissatisfied	Disreputable	$\rightarrow$	Increase
Dissatisfied	Unsure	$\rightarrow$	No change

[3] $r_s^d - \bar{r}_s^A \approx 1.28 > \sigma^s \approx 0.86$

After the reputation of each advisor is updated, the partition of buyers (R,DR,?) is updated using the criteria of (1).

5.4 Advisor cache

The purpose of the advisor cache is to store all the information that a buyer needs to maintain about a the advisors. For each advisor a the cache contains their reputation r_a^b, the reputation error σ^a, and a list of the seller's a has experience with.

When a buyer decides to ask about a set of potential sellers, the advisor cache is scanned to a set of advisors who have experience with some subset of the potential sellers. The disreputable advisors is removed and the set can be paired down even further by selecting only the advisors who have a sufficiently low reputation error.

6 Discussion

6.1 Cheap pseudonyms

Our approach as presented thus far is vunerable to sellers with bad reputations re-entering the market using pseudonyms, since the reputation assigned to new sellers is significantly higher than the lowest reputation. Friedman and Resnick present a some possible solutions that can be implemented on in addition to our approach to address the problem of cheap pseudonyms [10].

One solution offered is the use of a Trusted Third Party or TTP to provide what they call *once-in-a-lifetime identifiers*. The TTP requires some proof of the users real world identity before it issues the once-in-a-lifetime identifier. This identifier can be verified as being issued by the TTP, but contains no link to the real identify of the user. The mapping is between the user's real identity and identity within in the market is a secret kept by the TTP.

Another solution offered makes use of payments. A new user to the market must pay an entry fee which is distributed to all the users who are already a part of the market. In this way a user is provided with an ongoing incentive for staying in the market and a disincentive for re-entering under a different identity.

The use of either of these approaches while increasing the complexity of our system would not conflict with any of the methods proposed for avoiding deceptive sellers.

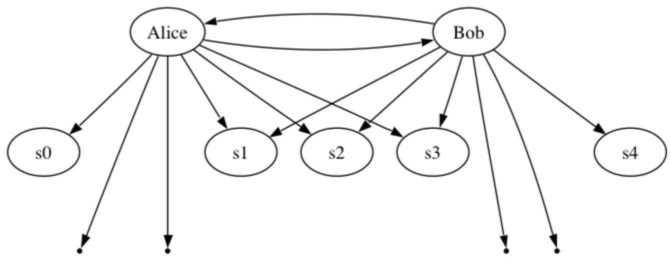

Figure 2: Alice and Bob and their experience

6.2 Cooperation

The influence of mechanism design and game theory has motivated our focus on the deception possible by sellers acting out of self-interest. It is beneficial for buyers to cooperate in most situations however, this understanding of group welfare is not something we can expect from our self-interested buyers in game theoretic models. Instead it is the *rules of interaction* provided by the mechanism designer that should provide the incentive for cooperation.

This paper provides only a basic foundation for motivating the cooperation of other buyers by identifying and minimizing the impact of deceptive buyers. This foundation could be extended to include notions of reciprocity which determine when a buyer would decide to cooperate and share seller reputation. The centralized matchmaker server could be used to pair up buyers whose cooperation would directly benefit each other. The advisor cache could be used to store information about which advisors helped the buyer in the past, indicating that the buyer should help later on. Future work will explore how we can provide incentive for cooperation by allowing agents to form small groups where agents are required to help each other.

There is a wealth of interesting problems associated with providing incentive for cooperation in reputation mechanisms of which most of the work cited in this paper barely scratches the surface. Most of the models with buyer communication which allow buyers to be self-interested enough to lie, do not address the situation where a buyer receives help from others and simply provides no help in return [7, 14, 17, 18, 19].

7 Related work

7.1 Yu and Singh

Yu and Singh present a similar approach to trust in multi-agent systems [17, 18]. An agent builds a reputation about 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 they 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.

In their most recent work [18] they use a different representation of reputation which is based on Dempster-Shaffer theory of evidence. 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, it does suffer from one problem which our model does not. In the model of Yu and Singh, the evidence for or against trustworthiness is not weighted by the value of a transaction. Thus after being happy with a beanie-baby that was bought from a correspondent and terribly unhappy that the truck I also bought from the correspondent was never delivered, I am left with reputation that lists one piece of evidence for and one against. In the model of Cohen and Tran adopted in this paper the model that correspondents reputation would go up slightly with the successful sale of the beanie-baby and would drop precipitously after having not delivered the truck.

Witnesses are sought out and correspondents evaluated one at a time. Our model takes advantage of situations in which there may be multiple sellers offering a good. By asking about a set of sellers, the buyer is free to chose among the sellers who have been deemed reputable based on some other criteria such as estimated value.

7.2 Sabater and Sierra

Sabater and Sierra developed REGRET [14], a model of reputation that takes into account the personal, social and ontological aspects of reputation. The ontological structure of reputation explored by Sabater and Sierra in which different aspects of reputation can be combined into a more complex reputation goes well beyond what we have addressed in this paper, but the way in which they address the personal and social aspects are similar to what we have done.

In their model, the direct interactions between a buyer and a seller would build the buyers personal reputation of the seller. This personal reputation captures some phenomena that ours does not, such as the relevance of current interactions compared to those far in the past. The social aspect of reputation includes a method to combine reputations of a seller held by multiple buyers using weights to adjust for the differences between buyers. While REGRET provides a robust method for modeling many aspects of a seller, it does not address the possibility of dishonesty among agents sharing reputation information in the social aspect. In comparison our model offers a method to identify dishonest advisors in some cases even before making a purchase.

7.3 Zacharia, Moukas and Maes

Zacharia, Moukas and Maes have proposed a collaborative reputation mechanism called Sporos for electronic marketplaces [19]. Sporos 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 weighs each rating by the reputation by finding the product of the two values.

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. If a new user cannot be distinguished from users who have been deceitful, it is unclear how a new user would convince enough other users to interact with it and build its reputation.

The paper while presenting another novel model for using chains of trust in highly connected networks, does not address *when* other agents should be consult, *what* criteria should be used to find these other agents, or *how* they should address the subjectivity in each agents ratings.

8 Conclusions

The focus of this paper was to provide the potential for optimal market outcomes by reducing behavior which is detrimental to the welfare of the market as a whole. Specifically we examined how a system can be designed to limit the effect of deceptive sellers (and buyers) from the perspective of a single agent in a multi-agent system.

The system developed allows the buyer to use indirect reputation gather from other buyers acting as advisors to judge the reputation of sellers for which there is no direct reputation information. This model assumes that the indirect reputations provided by advisors is subjective and may not be truthful.

This paper offers two approaches to addressing deceptive advisors. The buyer will model the reputation of advisors and only listen to those who are deemed reputable. Also a buyer will ignore advisors whose predictions are significantly different from their peers in order to reduce the impact of deceptive advisors when combining indirect reputation gathered from a group of advisors. Since this process happens before a seller is selected, a deceptive advisor could be detected and dealt with before that advisor has had the opportunity to fool the buyer even once.

We have studied the challenges arising from subjectivity in reputations that are shared between buyers and we have offered an approach to identify any systematic bias in the reputations of sellers common to a pair of buyers and interpret future seller reputations to correct for this bias.

By providing approaches to reduce the problems associated with subjectivity and the possibility of deception with advisors, this paper provides a solid foundation for use of indirect reputation to promote better market outcomes by reducing the impact of deceptive sellers.

9 Future Work

While this paper presents a novel approach for using reputation, it offers no formal evaluation of that model. Future work will include an empirical analysis of an implementation of a multi-agent system using our approach. This analysis will attempt to verify the hypothesis that the use of other buyers as advisors is beneficial when a buyer lacks reputation information about potential sellers. We can use the model provided by Cohen and Tran [6] as a base case since our approach reduces to their model when buyers do not communicate. To measure how far we have come towards meeting the challenges set out in section 3 the implementation should provide scenarios involving

- randomly assigned deceptive sellers and buyers drawn from different distributions
- buyers with different levels of subjectivity regarding their reputations
- deceptive sellers who attempt to shed bad reputations by re-entering the market as new users
- increasing numbers of buyers, sellers and transactions to measure how the system scales

We can assess many of our design decisions by implementing them in isolation and observing their impact on the system as a whole. It will be useful to investigate how the implementation reacts to the adjustment of model parameters towards finding optimal values. It will interesting to measure communication overhead to assess whether any benefits obtained through our approach outweigh the resulting increase in network load.

Future work will also focus on providing incentive for cooperation in conjunction with our approach. Mechanisms where agents are encouraged to help those who have helped them in the past is a good starting point. However, there are more complex scenarios in which there are an asymmetry in groups of agents who could benefit from cooperation. An example of this is a buyer a who can provide advice to agent b, who provides advice to agent c, who provides advice to agent a. Each agent is both giving and receiving, however since there is no direct symmetry from the perspective of a single agent it is difficult to know whether a cycle is formed or not. If no cycle is formed, than one or more agents may be getting a free ride by taking without giving.

diagram

Incorporating models of cooperation into multi-agent systems using reputation will lead us closer to fully realizing our assumptions of self-interested agents interacting with rules of interaction we provide to achieve some common goal.

10 Acknowledgements

I would like to thank Robin Cohen for her help inspiring and refining many of the ideas that went into this paper.

References

[1] Amazon.com auctions, auctions.amazon.com.

[2] ebay, www.ebay.com.

[3] E. Adar and B. Huberman. E. adar and b. huberman. free riding on gnutella. Technical report, Xerox PARC, 2002.

[4] George A Akerlof. The market for 'lemons': Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84(3):488–500, Aug. 1970. available at http://ideas.repec.org/a/tpr/qjecon/v84y1970i3p488-500.html.

[5] Silvia Breban and Julita Vassileva. Using inter-agent trust relationships for efficient coalition formation. In *Proceedings of the 15th Conference of the Canadian Society for Computational Studies of Intelligence on Advances in Artificial Intelligence*, pages 221–236. Springer-Verlag, 2002.

[6] Robin Cohen and Thomas Tran. A framework for modelling reputation to protect from infinite harm. 2004.

[7] Chrysanthos Dellarocas. Analyzing the economic efficiency of ebay-like online reputation reporting mechanisms. In *EC '01: Proceedings of the 3rd ACM conference on Electronic Commerce*, pages 171–179. ACM Press, 2001.

[8] Chysanthos Dellarocas. Immunizing online repuation reporting systems against unfair rating and discriminatory behavior. *E-Commerce*, October 2000.

[9] Michal Feldman, Kevin Lai, Ion Stoica, and John Chuang. Robust incentive techniques for peer-to-peer networks. In *EC '04: Proceedings of the 5th ACM conference on Electronic commerce*, pages 102–111. ACM Press, 2004.

[10] Eric J. Friedman and Paul Resnick. The social cost of cheap pseudonyms. 1999.

[11] Radu Jurca and Boi Faltings. "confess": Eliciting honest feedback without independent verification authorities. *Sixth International Workshop on Agent Mediated Electronic Commerce (AMEC VI)*, 2004.

[12] Sergio Marti and Hector Garcia-Molina. Limited reputation sharing in p2p systems. In *EC '04: Proceedings of the 5th ACM conference on Electronic commerce*, pages 91–101. ACM Press, 2004.

[13] Lik Mui, Mojdeh Mohtashemi, and Ari Halberstadt. Notions of reputation in multi-agents systems: a review. In *Proceedings of the first international joint conference on Autonomous agents and multiagent systems*, pages 280–287. ACM Press, 2002.

[14] Jordi Sabater and Carles Sierra. Regret: reputation in gregarious societies. In *AGENTS '01: Proceedings of the fifth international conference on Autonomous agents*, pages 194–195. ACM Press, 2001.

[15] R. G. Smith. The contract net protocol: high-level communication and control in a distributed problem solver. pages 357–366, 1988.

[16] Hal R. Varian. Economic mechanism design for computerized agents. In *Proceedings of the First USENIX Workshop on Electronic Commerce*, July 1995.

[17] Bin Yu and Munindar P. 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. Springer-Verlag, 2000.

[18] Bin Yu and Munindar P. Singh. Detecting deception in reputation management. In *Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, pages 73–80. ACM Press, 2003.

[19] Giorgos Zacharia, Alexandros Moukas, and Pattie Maes. Collaborative reputation mechanisms in electronic marketplaces. In *32nd Hawaii International Conference on System Sciences*, 1999.