

Bayesian Reputation Modeling in E-Marketplaces Sensitive to Subjectivity, Deception and Change

Kevin Regan and Pascal Poupart and Robin Cohen

David R. Cheriton School of Computer Science; University of Waterloo, Waterloo, Ontario, Canada
{kmregan, ppoupart, rcohen}@uwaterloo.ca

Abstract

We present a model for buying agents in e-marketplaces to interpret evaluations of sellers provided by other buying agents, known as advisors. The interpretation of seller evaluations is complicated by the inherent subjectivity of each advisor, the possibility that advisors may deliberately provide misleading evaluations to deceive competitors and the dynamic nature of seller and advisor behaviours that may naturally change seller evaluations over time. Using a Bayesian approach, we demonstrate how to cope with subjectivity, deception and change in a principled way. More specifically, by modeling seller properties and advisor evaluation functions as dynamic random variables, buyers can progressively learn a probabilistic model that naturally and “correctly” calibrates the interpretation of seller evaluations without having to resort to heuristics to explicitly detect and filter/discount unreliable seller evaluations. Our model, called BLADE, is shown empirically to achieve lower mean error in the estimation of seller properties when compared to other models for reasoning about advisor ratings of sellers in electronic marketplaces.

1 Introduction

Consider an electronic marketplace where buyers who are interested in acquiring goods are represented by buying agents, running algorithms designed to make effective selection of selling agents. In scenarios where buyers can only inspect the goods after purchase, it is important to develop strategies to determine the reliability of the sellers prior to a purchase. Various models have been developed to enable a buying agent to learn from past experience with sellers, in order to make effective decisions in future transactions (Tran & Cohen 2004; Vidal & Durfee 1996). Another approach is to enable a buying agent to obtain advice from other buying agents in the marketplace (advisors) (e.g. (Zacharia, Moukas, & Maes 1999; Whitby, Josang, & Indulska 2004; Yu & Singh 2003; Teacy *et al.* 2005)).

The challenge then becomes how to interpret information about sellers from advisor agents according to: subjective differences between the buyer and advisor, possible deception on the part of the advisor and changes over time in views of advisors and performance of sellers.

A common approach is to have advisors report a single reputation rating for each seller to the buying agent and then to provide a mechanism for detecting and addressing any unreliable ratings that are provided. Detection is usually done by verifying whether the ratings of an advisor differ significantly from some aggregate statistic of the ratings of all the advisors together (Whitby, Josang, & Indulska 2004) or from the buyer’s own experience (Teacy *et al.* 2005; Yu & Singh 2003). Ratings deemed unreliable are then discounted in the aggregation or simply filtered out. While this is a reasonable approach, a significant amount of information may be lost as a result of discounting/filtering unreliable ratings.

In some cases, when an advisor is known to consistently bias its ratings (e.g. always exaggerating positively or negatively, or always reporting the opposite of what it thinks), it is in fact possible to “re-interpret” the ratings and to reduce the need for discounting. We propose a new model, called BLADE (Bayesian Learning to Adapt to Deception in E-Marketplaces), that re-interprets unreliable ratings whenever possible. More specifically, we develop a Bayesian approach that models both a set of seller properties and the ratings provided by the advisors, considering each to be random variables. Then, the evaluation function used by each advisor in reporting a rating to a buying agent is learned over time and is used by our system to re-interpret ratings (in spite of any subjective differences or possible deception) in terms of those properties of the seller that matter to the buyer.

We demonstrate empirically that BLADE achieves higher accuracy in learning about sellers than two existing Bayesian models (BRS (Whitby, Josang, & Indulska 2004) and TRAVOS (Teacy *et al.* 2005)) since it can often reduce the need for discounting/filtering (potentially unreliable) ratings provided by subjective and/or deceptive advisors. We also discuss how to adjust our model to cope with change in seller and advisor behaviour, over time and present experimental results to demonstrate the value of our approach. Our conclusion is that it is possible to apply a principled method of Bayesian learning to acquire deeper models of advisors and sellers that make more effective use of the information that is modeled. This is an important step in providing buyers with an effective method of selecting reputable sellers with whom to do business in electronic marketplaces.

2 Background

In order to appreciate the value of our approach to modeling reputation in electronic marketplaces, it is important to understand some of the competing social network models proposed by other researchers. For example, Histos, proposed by Zacharia et al. (1999) provides for chains of trust among agents, but does not have a mechanism to adjust for subjective differences or to deal with deceptive advice. Sabater and Sierra (2001) offer a rich multi-faceted model of reputation, but do not explicitly account for dishonesty among agents. Furthermore, these approaches are not probabilistic, and as such operate using various formulae with weights or thresholds.

Recent work has led to more principled models grounded in Dempster-Schaefer theory and probability theory but still including heuristics for detecting and discounting unreliable ratings. For instance, Yu and Singh (2003) adopt Dempster-Schaefer theory to model the amount of evidence for and against a reputation value. Deception and subjectivity are addressed by taking a weighted majority of the ratings of each advisor according to how successful each advisor has been at predicting seller reputations. Whitby et al. (2004) and Teacy et al. (2005) adopt probability theory to develop two Bayesian models respectively called BRS and TRAVOS. Both models use ratings corresponding to the number of satisfactory and unsatisfactory seller interactions to construct a Beta distribution representing the seller’s reputation. To mitigate the effect of deceptive advice, TRAVOS computes a probability of accuracy measuring the likelihood that an advisor’s rating matches the buyer’s own opinion. A heuristic is then used to effectively discount ratings by a factor corresponding to this accuracy probability. In contrast, BRS uses the combined ratings of all advisors to eliminate the ratings that differ significantly from the majority, assuming that the majority is right.

While these models are grounded in principled theories, they still make use of heuristics to deal with unreliable ratings. By weighting, discounting or eliminating ratings, useful information may be thrown away. For instance, when an advisor offers consistent yet deceptive/subjective advice, it may be possible to adjust the rating instead of diminishing its importance. In the next section, we introduce our model, BLADE, and discuss how it interprets information from advisors about sellers in a more inclusive and effective manner.

3 BLADE model

We now describe our Bayesian modeling approach to handle subjectivity, deception and change, called BLADE. We first describe a Bayesian network (BN) that enables buyers to learn seller properties and advisor evaluation functions. Then, this basic BN is expanded into a dynamic Bayesian network (DBN) in Section 3.2 to allow seller and advisor behaviours to change over time. Throughout the paper we use the convention that capital letters (e.g. F) denote variables, lowercase letters denote values (e.g. f), bold letters denote vectors or sets (e.g. $\mathbf{F} = \{F_1, \dots, F_k\}$).

We begin by assuming that the value derived from a purchase is a function of multiple aspects of the transaction. For

example, a buyer’s utility¹ U^b may depend on the shipping time and the condition of the good purchased such as the precision of the machining of automotive parts in a car manufacturing setting. In general, the aspects of the purchase that determine the buyer’s utility are determined by what can be considered some intrinsic properties of the seller. We denote by $\mathbf{F}^s = \{F_1^s, \dots, F_k^s\}$ the set of features (or properties) that seller s exhibits in a given transaction. For our purposes, each feature F_i^s can take on a finite number of discrete values often simply referred as f_i^s . For instance, a feature representing shipping time could take on the values $\{on\ time, one\ day\ late, one\ week\ late, more\ than\ one\ week\ late, did\ not\ arrive\}$. Since a seller’s behaviour is not deterministic (i.e. the values of the features may vary from one transaction to another), we regard seller features F_i^s as random variables whose values f_i^s are drawn from respective multinomial distributions $\Pr(F_i^s)$. Those multinomial distributions are the quantities that a buyer would like to estimate since they correspond to the intrinsic likelihood that each seller will exhibit some properties in a transaction.

Each buyer can gather information about the distribution from which the seller properties are drawn by purchasing goods from that seller, and by collecting information from other buyers in the marketplace who are acting as advisors. The information given by an advisor a about seller s takes the form of a rating r_s^a chosen from a finite number of discrete values such as $\{satisfied, unsatisfied\}$ or a number on a scale of 1 to 10 for example. Since ratings may vary with each transaction, R_s^a is considered a random variable. We allow the domain of each R_s^a to be different, reflecting the fact that different advisors may choose a different range of values to report their ratings. Note also that we do not assume that identical ratings necessarily mean the same thing when given by different advisors nor that the ratings necessarily reflect the advisor’s utility since an advisor may try to deceive by reporting misleading ratings. In general, we allow the evaluation function that maps the aspects of a purchase to the rating to be subjective and therefore specific to each advisor. Since this mapping may not always be a deterministic function of the seller properties due to inherent noise or the presence of additional factors not captured by the seller properties, the evaluation function is modeled as a conditional probability distribution $\Pr(R_s^a | \mathbf{F}^s)$. Figure 1 illustrates by a BN the probabilistic dependencies of buyer utility and advisor ratings on seller properties.

In scenarios where a seller has a large number of features that would yield an intractable BN, we can use feature selection techniques from machine learning (Blum & Langley 1997) to reduce the number of features. Note also that the BN in Figure 1 shows feature variables that are independent of each other simply to keep the exposition simple. Nothing in our model prevents the addition of arcs between the feature variables to introduce dependencies.

The inherent behaviour of each seller (i.e., $\Pr(F_i^s)$) and the evaluation function used by each advisor (i.e., $\Pr(R_s^a | \mathbf{F}^s)$) are not generally known by a buyer. Nevertheless, those distributions can be learned as a buyer inter-

¹Utility derived by a buyer from a transaction.




Figure 1: Bayesian Network illustrating the conditional dependence of a buyer utility U_s^b and three advisor ratings $R_s^{a1}, R_s^{a2}, R_s^{a3}$ on three features F_1^s, F_2^s, F_3^s of seller s .

acts with sellers and advisors. We use a Bayesian learning approach (Heckerman 1998), which allows a buyer to cope with deception in a principled way, that is, without having to resort to any heuristic to explicitly detect/filter misleading ratings. Bayesian learning proceeds by modeling each unknown with a new random variable θ . For instance, let $\theta_i^s = \Pr(F_i^s)$ encode the unknown properties for each feature i of seller s . Similarly, let $\theta_f^a = \Pr(R_s^a | \mathbf{f})$ encode the unknown multinomial over ratings given by advisor a when the seller exhibits feature values \mathbf{f} .

With these additional variables, we obtain the BN in Figure 2 for three advisors, one seller and three features per seller and the BN in Figure 3 for three advisors, two sellers and three features per seller. The seller behaviours and advisor evaluation functions are now conditioned on the θ variables allowing us to specify $\Pr(F_i^s | \theta_i^s) = \theta_i^s$ and $\Pr(R_s^a | \mathbf{f}, \theta_f^a) = \theta_f^a$. We also represent prior distributions over each θ with a Dirichlet distribution since Dirichlets are conjugate priors of multinomials (DeGroot 1970). A Dirichlet distribution $\mathcal{D}(\theta; \mathbf{n}) = k \prod_i \theta_i^{n_i - 1}$ over a multinomial θ is parameterized by positive numbers n_i (known as hyperparameters) such that $n_i - 1$ can be interpreted as the number of times that the θ_i -probability event has been observed. The expectation of a Dirichlet is given by a ratio of its hyperparameters (e.g., $E(\theta_i) = \int_{\theta} \theta_i \mathcal{D}(\theta; \mathbf{n}) d\theta = n_i / \sum_i n_i$).




Figure 2: Bayesian Network model for a buyer, three advisors, one seller and three features per seller. Note that θ^a denotes the set of θ_f^a variables for all \mathbf{f} .

Given the BN in Figure 2, the task of learning seller behaviours and advisor evaluation functions becomes a question of inference. Initially, a buyer starts with priors $\Pr(\theta_i^s)$ and $\Pr(\theta_f^a)$ encoding the likelihood of behaviours for each seller and evaluation functions for each advisor. As the buyer interacts with sellers and advisors, it gets to observe the seller properties of each transaction it enters and the ratings of the transactions that advisors report. The seller features \mathbf{f}^s observed are used as evidence to compute the




Figure 3: Bayesian Network model for a buyer, three advisors, two sellers and three features per seller. Note that θ^a denotes the set of θ_f^a variables for all \mathbf{f} .

posteriors $\Pr(\theta_i^s | \mathbf{f}^s)$ and $\Pr(\theta_f^a | \mathbf{f}^s)$. Similarly, the ratings r_s^a observed are used as evidence to compute the posteriors $\Pr(\theta_i^s | r_s^a)$ and $\Pr(\theta_f^a | r_s^a)$.

Even though we have a hybrid BN (R_s^a and F_i^s are discrete variables while θ_i^s and θ_f^a are continuous variables), those posteriors can be computed in closed form by exploiting the properties of Dirichlets. For instance, $\Pr(\theta_1^s | r_s^a)$ is computed in the usual way by multiplying all the BN conditional probability distributions and summing (or integrating) over all unobserved variables. We omit the seller index s in the derivation below to simplify the notation.

$$\Pr(\theta_i | r^a) = k \sum_{\mathbf{f}} \left[\int_{\theta^a} \Pr(\theta^a) \Pr(r^a | \mathbf{f}, \theta^a) d\theta^a \right] \times \left[\prod_{j \neq i} \int_{\theta_j} \Pr(\theta_j) \Pr(f_j | \theta_j) d\theta_j \right] \Pr(\theta_i) \Pr(f_i | \theta_i) \quad (1)$$

$$= k \sum_{\mathbf{f}} \left[\int_{\theta^a} \Pr(\theta^a) \theta_{r, \mathbf{f}}^a d\theta^a \right] \times \left[\prod_{j \neq i} \int_{\theta_j} \Pr(\theta_j) \theta_{j, f_j} d\theta_j \right] \Pr(\theta_i) \theta_{i, f_i} \quad (2)$$

$$= k \sum_{\mathbf{f}} E(\theta_{r, \mathbf{f}}^a) \left[\prod_{j \neq i} E(\theta_{j, f_j}) \right] \Pr(\theta_i) \theta_{i, f_i} \quad (3)$$

$$= \sum_{f_i} c_{f_i} \mathcal{D}_{f_i}(\theta_i) \quad (4)$$

Eq. 2 is obtained by substituting $\Pr(f_i | \theta_i)$ and $\Pr(r^a | \mathbf{f}, \theta^a)$ by θ_{i, f_i} and $\theta_{r, \mathbf{f}}^a$. Eq. 3 follows from the fact that each integral amounts to an expectation. Assuming $\Pr(\theta_i) = \mathcal{D}(\theta_i)$ then $\Pr(\theta_i) \theta_{i, f_i} = \mathcal{D}_{f_i}(\theta_i)$, which is the same Dirichlet as $\mathcal{D}(\theta_i)$, but with the hyperparameter for f_i incremented by 1. Since each expectation is a scalar, Eq. 4 shows that $\Pr(\theta_i | r^a)$ is simply a mixture of Dirichlets. While this derivation may look complex the computation of the final mixture is very simple since we only need to compute Dirichlet expectations that have a closed form consisting of a ratio of hyperparameters.

In general, when learning with a BN like the one in Figure 2, the posteriors of any θ given some evidence are always

mixtures of Dirichlets (Heckerman 1998). Hence we omit the derivations for $\Pr(\theta_i^s | \mathbf{f}^s)$, $\Pr(\theta_f^a | \mathbf{f}^s)$ and $\Pr(\theta_f^a | r_s^a)$ since they are similar.

As we incorporate seller properties observed in direct transactions and ratings reported by advisors, the distributions over θ^s and θ^a will be represented as a mixture of Dirichlets; however, the number of Dirichlets in each mixture may grow each time a new posterior is calculated. For instance in the case of $\Pr(\theta_i | r^a)$ the number of Dirichlets is multiplied by $|F_i|$ (i.e., size of the domain of F_i), which leads to an exponential growth. To avoid storing an exponential number of Dirichlets, we approximate mixtures of Dirichlets by a single Dirichlet whose hyper-parameters are a mixture of the hyper-parameters of each Dirichlet.

$$\sum_i c_i \mathcal{D}(\theta; n_{i1}, \dots, n_{ik}) \approx \mathcal{D}(\theta; \sum_i c_i n_{i1}, \dots, \sum_i c_i n_{ik}) \quad (5)$$

Note however that this approximation has the benefit of preserving expectations (Cowell 1998):

$$\begin{aligned} \int_{\theta} \theta_j \sum_i c_i \mathcal{D}(\theta; n_{i1}, \dots, n_{ik}) d\theta &= \sum_i c_i n_{ij} / \sum_i c_i n_{ij} \\ &= \int_{\theta} \theta_j \mathcal{D}(\theta; \sum_i c_i n_{i1}, \dots, \sum_i c_i n_{ik}) d\theta \end{aligned}$$

Since all the calculations we make only use the expectation of the θ variables, this approximation is not really an approximation as it preserves a sufficient statistic that allows us to do calculations exactly.

Let’s now examine how the Bayesian learning approach copes with deception. In our model, advisor ratings and seller properties are random variables. The greater the correlation between them the higher their mutual information and the easier it is to infer the value of one variable given the value of the other variable. An advisor evaluation function $\Pr(R_s^a | \mathbf{F}^s)$ essentially encodes the correlations between ratings and seller properties. The more deterministic the evaluation function learned is, the stronger the correlations will be. So if an advisor tends to report a unique rating (the actual rating doesn’t matter, as long as it is unique) for each value of the seller features, then it is possible to infer back the seller properties. For instance if an advisor tends to report *satisfied* for one seller when the buyer has established that the seller usually delivers *late* and *unsatisfied* for another seller when the buyer has established that this seller usually delivers *on time*, our approach will exploit the correlation between *late* and *satisfied* to allow the buyer to interpret *satisfied* from this advisor as meaning *late*. Hence a buyer can equally make use of ratings from honest and dishonest advisors as long as they are consistent (i.e. use a fairly deterministic evaluation function).

If an advisor’s evaluation function is quite stochastic and in the extreme completely random, then it will be more difficult to infer seller properties. Note that a buyer is not adversely affected when ratings are weakly correlated with the seller properties, since the algorithm doesn’t try to infer a single seller property for a given rating, but rather a distribution over seller properties. Hence, when an advisor gives random ratings, the algorithm will simply infer that all seller properties are equally likely which amounts to ignoring the

rating. When the ratings are weakly correlated then the algorithm will infer that some properties are only slightly more likely than others, which can be thought as providing only a little bit of information. In general, Bayesian learning extracts just the “right” amount of information from each rating based on the amount of correlation between ratings and seller properties.

We can imagine some cases in which an advisor colludes with a seller, deviating from its standard evaluation function to offer inflated ratings for that particular seller. The advisor could most effectively collude by reporting ratings that are highly correlated with seller properties for every seller except the colluding seller. However, it is reasonable to assume that the ratings of many advisors will be used to model each seller’s properties and thus a colluding advisor will have a minimal impact.

3.1 Example

Let us return to the scenario involving sellers with a single seller property F^s which can take the value *late* or *on time* and a single advisor who will offer ratings of either *unsatisfied* or *satisfied*. This example will walk through how, given some information about a seller s_1 , a buyer can learn the evaluation function of the advisor a and use this information to infer the property of an unknown seller s_2 . To aid in conceptualizing the state of the parameters θ^{s_1} , θ^{s_2} , θ_1^a and θ_2^a , we visualize the distribution over each θ using the probability density function.

We begin by assuming that our buyer has already had 10 interactions with s_1 where the delivered good was *on time* 9 times and *late* only once. The buyer calculates a posterior $\Pr(\theta^{s_1} | \mathbf{f}^s)$ after each interaction learning that the seller s_1 is usually *on time* (with a probability density concentrated near 0.9). Note that in scenarios in which the buyer has had no experience with advisors (i.e. a new buyer entering the market), our probabilistic framework allows for an informative prior to model the advisor’s evaluation function. We leave the generation of such a prior to future work, but note that it serves as a starting point. Given any prior, our technique for adjusting for change discussed in the next section will eventually converge on a good representation for the advisor evaluation function. Now, at this point, our buyer has




Figure 4: Parameters before gathering advisor ratings

no information about the advisor a or the seller s_2 (Fig 4). The buyer gathers a set of ratings about the known seller s_1 from the advisor; the advisor reports that it is usually *unsatisfied* providing 9 *unsatisfied* ratings and 1 *satisfied* rating. The ratings r provided by the advisor are used to calculate the posterior $\Pr(\theta_1^a | r)$ and $\Pr(\theta_2^a | r)$ and since our buyer has already learned that the seller s_1 is usually on time, it has

effectively learned that there is a correlation between the rating of *unsatisfied* from this advisor with the seller property *on time*. Given a set of ratings from the advisor about an




Figure 5: Parameters after ratings for known seller

unknown seller s_2 , our buyer can interpret these ratings to infer back the seller properties which led to each rating. The advisor says that it is very happy with s_2 being *satisfied* 9 times and only *unsatisfied* once, but because the buyer has learned the advisor’s evaluation function, the buyer is able to interpret being mostly *satisfied* as indicating that the seller s_2 is usually late. Note that it is possible that the advisor




Figure 6: Parameters after ratings for unknown seller

is being deliberately deceptive - e.g. it in fact values sellers who deliver on time but reports the opposite rating, when asked. It is also possible that there is just a subjective difference - e.g. the advisor is satisfied when the good arrives in excellent condition (a different seller property) and this correlates with arriving late. Without explicitly knowing the seller property determining the advisor ratings, BLADE is able to establish a correlation between the ratings reported and known seller properties and use this to learn the advisor’s subjective evaluation function, to infer the seller properties of unknown sellers.

In contrast with this example, note that subjective/deceptive advisors are not expected to report ratings that are completely anti-correlated with seller properties. Our approach is equally capable of re-interpreting inaccurate ratings from such advisors as long as their evaluation function consistently maps each seller property to a unique rating.

Note that the example illustrated in this section is one in which the buyer is able to learn a strong correlation between the advisor reporting *unsatisfied* and the seller being *on time*. If the advisor has responses that are more weakly correlated (for example, reporting *unsatisfied* 7 out of 10 times), then we will simply be learning less information about the seller properties.

3.2 Dynamic sellers and advisors

The behaviour of sellers and advisors may shift over time. For example a selling agent providing automotive parts may improve the property representing the precision of a particular dimension of the delivered part after an upgrade of the

seller’s machining process. Similarly, the evaluation function of an advisor may change as a result of a change in its underlying utility function or perhaps simply because it feels like suddenly misleading.

As is, BLADE weighs equally each piece of evidence (e.g. rating or observed seller property). Thus, if the buyer has learned a seller property with high certainty and the property changes, a significant amount of new evidence would be necessary to change the buyer’s belief about this seller property. To effectively adjust what the buyer has learned to account for change, we incorporate the intuition that recent evidence is more representative of a seller property or buyer utility function, than what has been observed in the past.

The hyperparameters $n_1 \dots n_k$ of a Dirichlet distribution $\mathcal{D}(\theta; n_1 \dots n_k)$ can be thought of as representing the number of past occurrences of each θ_i -event. To give more weight to recent observations, we use the same technique as BRS (Whitby, Josang, & Indulska 2004): we construct a new Dirichlet \mathcal{D}' after each observation and calculation of posterior by scaling down each hyperparameter by a constant factor $\delta \in [0, 1]$ as follows: $\mathcal{D}'(\theta; \delta n_1, \dots, \delta n_k)$. This approach increases the uncertainty in the belief distribution represented by the Dirichlet, giving more weight to subsequent observations and allowing the distribution to change more quickly to accommodate a change in the pattern of evidence. An experiment in the next section demonstrates the impact of δ on the ability of a buyer to adapt to changes in the seller properties.

Figure 6 illustrates a Dynamic Bayesian Network representation of this process. At each time step the hyperparameters of the Dirichlet for each θ are calculated by scaling by δ . The variables R and F are conditionally independent of their values from previous time steps given all θ s.

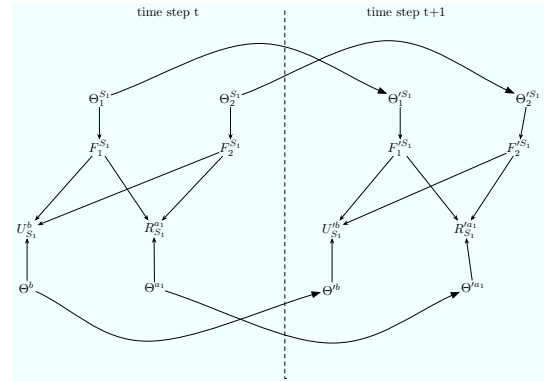


Figure 7: Dynamic Bayesian Network model incorporating change for a buyer, one advisor, one seller and 2 features per seller.

4 Experiments

We evaluate our approach using a simulated market scenario to examine how effectively we are able cope with subjectivity and deception by learning the evaluation functions of




Figure 8: (a) Deception experiment (b) Noise experiment (c) Seller change experiment

advisor agents. We compare our approach to implementations of TRAVOS and BRS, which both aggregate advisor information to construct a single seller reputation while mitigating the effect of deceptive advice. Both systems collect *opinions* from the advisors which consist of the number of satisfactory and unsatisfactory interactions the advisor has had with a seller.

To ensure a reasonable comparison, the sellers in our market have one seller property which can be thought of as the reputation of that seller. The advisors use one of two possible evaluation functions: *truthful* advisors will directly report the seller property they observed (essentially indicating whether they were satisfied or unsatisfied), while *deceptive* advisors will report the opposite of the seller property they observed. For the purposes of our evaluation, the market consists of a set of 11 *known* sellers that we interact with directly and 11 *unknown* sellers that we only gain information about through a set of 20 advisors. Each seller is assigned an intrinsic reputation $\Pr(F^s = \textit{satisfied})$ from the set $\{0, 0.1, 0.2, \dots, 0.9, 1\}$. Our experiment varies the percentage of *deceptive* advisors and assesses how well each system is able to integrate the information given by these deceptive advisors. In the first step of our evaluation the buyer interacts with the set of *known* sellers and records the outcome. Then, the buyer collects a set of opinions from each advisor about a randomly drawn subset of the *known* set of sellers, where each opinion represents 20 advisor-seller interactions randomly determined by the seller’s reputation $\Pr(F^s = \textit{satisfied})$. The information in this first step is used by BRS and TRAVOS to detect deceptive advisors and in BLADE to estimate the advisor evaluation function. Next, our buyer collects opinions about the *unknown* sellers in the market (once again based on 20 advisor-seller interactions) and estimates the reputation value e_s , which is the expected reputation of each *unknown* seller. Each of the three systems is evaluated by computing the *mean error* $= \frac{1}{N} \sum_{s=1}^N |\Pr(F^s = \textit{satisfied}) - e_s|$ over all the $N = 11$ unknown selling agents s .

Figure 7a graphs change in the *mean error* of each system as the percentage of deceptive advisors grows. The *mean error* of the reputation estimates constructed by BRS and TRAVOS increases as the percentage of deceptive advisors

grows. Recall that both systems discount or eliminate opinions which do not match previous opinions (of the group of advisors or the buyer’s own experience), and as the amount of deceptive advisors increases the amount of information useful to each system decreases, increasing the *mean error*. Our system assumes that the advisor opinions are derived from an evaluation function which may not be a direct mapping of the seller properties experienced by the advisor. The result is that we are able to correctly interpret the opinions of the deceptive advisors and for our system the percentage of deceptive agents has no impact on the *mean error* of the reputation estimate.

The same experimental setup was used to examine how inconsistent advisors affected BLADE and the other models. Instead of deceptive advisors we introduce noisy advisors. Inconsistent advisors are modelled by adding Gaussian noise with a standard deviation of 0.75 to the distribution generating the set of ratings provided by the noisy advisor. We then varied the percentage of noisy advisors in the market and the results are illustrated in Figure 7b. As expected, the mean error increases as the inconsistency due to the percentage of noisy advisors increases, but BLADE’s performance is on par with the other models².

To verify the effectiveness of our mechanism for incorporating change we focus on seller properties, noting that the approach to change of advisor evaluation functions is identical. We simulate a simple scenario in which a sequence of ratings are reported by an advisor and at the midpoint of this sequence the actual seller property generating the ratings is changed. We compare how well the buyer is able to initially learn the seller property and adjust what it has learned after the actual seller property has changed.

To focus on how a change in the distribution over actual seller properties affects our buyer’s learning of these properties, we assume that the buyer has already learned the advisor’s evaluation function with high certainty. In our evaluation we learned two binary seller properties F_1^s and F_2^s

²From an information theoretic perspective, there is no information in completely random ratings, so it is normal that none of the techniques are doing well. It makes sense to ignore random ratings, which is what BLADE does implicitly and BRS/TRAVOS do explicitly. This is why all techniques perform equally well.

whose actual distributions were $\Pr(F_1^s = 1) = 0.7$ and $\Pr(F_2^s = 1) = 0.3$. During the first half of the sequence our advisor would report ratings derived from observing seller properties drawn from the actual distribution. Half way through the sequence, the actual distribution of seller properties is shifted to $\Pr(F_1^s = 1) = 0.9$ and $\Pr(F_2^s = 1) = 0.1$. After each reported rating we record the expectation of the parameters representing seller properties θ_1^s and θ_2^s and compare this with the actual seller property distributions to generate the $error = E[\theta_1^s] - \Pr(F_1^s)$; we then average this error over 50 runs. This evaluation was run using our standard model without any adjustment for change and repeated after incorporating our methods to adjust for change. Figure 7c graphs the mean error at each time step when a rating is reported. We can see that when the actual seller properties are shifted, the mean error immediately goes up, but the rate at which this mean error then drops is far greater when change is incorporated. Also note, that the initial learning of the seller properties is not as stable when we incorporate possible seller change. This is due to the scaling down of the hyper-parameters specifying the distribution over θ_1^s and θ_2^s which has the dual effect of allowing for change to be more quickly incorporated, and increasing uncertainty.

5 Discussion, Conclusions & Future Work

It is useful to note that both BRS and TRAVOS can be recast as Bayesian Networks of the form $\theta \rightarrow F_{rep}^s \rightarrow R^a$ in which there is a single seller feature F_{rep}^s corresponding to its reputation. Estimating the reputability of a seller is achieved by computing the posterior $\Pr(\theta|r^a)$ based on ratings provided by advisors. Here, the evaluation function $\Pr(R^a|F_{rep}^s)$ is the identity function, assuming that ratings directly reflect the reputation of the seller. Viewed in this light, it is then clear that our model offers a generalization of BRS and TRAVOS with two significant improvements. First we go beyond a single reputation value and model a set of seller properties, allowing for a buyer to calculate the expected utility of a purchase, and not just the level of trust. Secondly, we use a second parameter θ^a to model and learn the advisor's evaluation function $\Pr(R^a|f_{rep}^s)$, which allows the buyer to re-interpret any ratings that are not a direct mapping of the seller properties. This re-interpretation allows us to infer back the seller properties despite subjective or deceptive ratings, as long as they are consistent with the evaluation function we have learned. Cases where there are merely subjective differences in evaluating sellers between advisor and buyer are in fact quite reasonable to expect in any electronic marketplace, and yet previous research has focused primarily on detecting and addressing deception.

The BLADE model provides a principled method of Bayesian learning in order to acquire deeper models of advisors and sellers that use the information available more effectively without relying on discounting or filtering. In our approach, then, both subjective differences and deception are addressed by the same mechanism; we are, moreover, able to incorporate methods for reasoning about change that favour more recent information. The result is an improved methodology for modeling trust between agents in electronic

marketplaces, as an important part of a buying agent's decision making about purchases, to avoid disreputable sellers.

One promising area of future work involves extending our model with actions and rewards within a decision theoretic framework to actively select which advisors to consult and which sellers to inquire about in order to minimize the burden of information collection. As a first step, we would explore the use of a POMDP framework similar to the one proposed in (Regan, Cohen, & Poupart 2005) to generate policies for selecting optimal actions. Our model could also be enriched by using iterative preference elicitation to learn the buyer utility function (Boutilier *et al.* 2005).

Acknowledgements

The authors gratefully acknowledge the financial support of Canada's Natural Science and Engineering Research Council (NSERC).

References

- Blum, A. L., and Langley, P. 1997. Selection of relevant features and examples in machine learning. *Artif. Intell.* 97(1-2):245–271.
- Boutilier, C.; Patrascu, R.; Poupart, P.; and Schuurmans, D. 2005. Regret-based utility elicitation in constraint-based decision problems. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 929–934.
- Cowell, R. G. 1998. Mixture reduction via predictive scores. *Statistics and Computing* 8:97–103.
- DeGroot, M. H. 1970. *Optimal Statistical Decisions*. New York: McGraw-Hill.
- Heckerman, D. 1998. A tutorial on learning with Bayesian networks. In Jordan, M. I., ed., *Learning in Graphical Models*. Kluwer.
- Regan, K.; Cohen, R.; and Poupart, P. 2005. The Advisor-POMDP: A principled approach to trust through reputation in electronic markets. In *PST*, 121–130.
- Sabater, J., and Sierra, C. 2001. Regret: reputation in gregarious societies. In *Autonomous Agents*, 194–195.
- Teacy, W. T. L.; Patel, J.; Jennings, N. R.; and Luck, M. 2005. Coping with inaccurate reputation sources: experimental analysis of a probabilistic trust model. In *AAMAS*, 997–1004.
- Tran, T., and Cohen, R. 2004. Improving user satisfaction in agent-based electronic marketplaces by reputation modelling and adjustable product quality. In *AAMAS*, 828–835.
- Vidal, J. M., and Durfee, E. H. 1996. The impact of nested agent models in an information economy. In *Second International Conference on Multi-Agent Systems*, 377–384.
- Whitby, A.; Josang, A.; and Indulska, J. 2004. Filtering out unfair ratings in bayesian reputation systems. In *Intl. Workshop on Trust in Agent Societies*.
- Yu, B., and Singh, M. P. 2003. Detecting deception in reputation management. In *AAMAS*, 73–80.
- Zacharia, G.; Moukas, A.; and Maes, P. 1999. Collaborative reputation mechanisms in electronic marketplaces. In *Hawaii International Conference on System Science*.