



Evaluating the trustworthiness of advice about seller agents in e-marketplaces: A personalized approach

Jie Zhang*, Robin Cohen

David R. Cheriton School of Computer Science, University of Waterloo, Waterloo, ON, Canada N2L 3G1

ARTICLE INFO

Article history:

Received 3 March 2008

Received in revised form 3 March 2008

Accepted 3 March 2008

Available online 13 March 2008

Keywords:

Trust

Buyer and seller agents

Private and public reputation

Unfair ratings

Electronic marketplaces

ABSTRACT

In this paper, we present a model for evaluating the trustworthiness of advice about seller agents in electronic marketplaces. In particular, we propose a novel personalized approach for effectively handling unfair ratings of sellers provided to buyer agents from other buyers (called advisors). Our approach offers flexibility for buyers to weight their value for private and public knowledge about advisors. A personalized approach is proposed as well for buyers to model the trustworthiness of sellers, based on the advice provided. Experimental results demonstrate that our approach can effectively model trustworthiness for both advisors and sellers, even when there are large numbers of unfair ratings.

© 2008 Elsevier B.V. All rights reserved.

1. Introduction

As the enterprise of electronic commerce becomes increasingly popular, worldwide, one challenge that arises is to ensure that organizations participating in e-commerce have sufficient trust in order to bring their businesses on-line. In order to assist both individual buyers and business organizations in conducting both B2B and B2C e-commerce, researchers in artificial intelligence have been designing intelligent agents to perform the tasks of buying or selling, on behalf of their human clients. While these agents assist in offloading the processing required by people in order to find the best business partnerships, it then becomes critical for these agents to make effective decisions, in order to engender the trust of their users.

In this paper, we examine one particular problem that arises when buyer agents elicit opinions about seller agents from other buyer agents in the marketplace: the issue of possible unfair ratings. To explain, in an agent-based electronic marketplace, agents are self-interested. They interact with each other to achieve their own goals. Seller agents sell products to buyer agents and try to maximize their profit. Buyer agents try to gain good products in terms of, for example, high quality and low prices. To ensure good interactions amongst agents, a trust and reputation mechanism

provides important social control in electronic marketplaces. In such a system, agents can rate each other. Agents estimate each other's trustworthiness according to those ratings and choose the most trustworthy ones to interact with. However, as buyers seek to find trustworthy sellers, they may be deceived by unfair ratings of sellers provided by other buyer agents, for their personal gain. Dellarocas [3] distinguishes unfair ratings as unfairly high ratings and unfairly low ratings. Unfairly high ratings may be used to increase seller agents' reputations. They are often referred as "ballot stuffing". Unfairly low ratings of a seller agent may be provided by buyer agents that cooperate with other seller agents to drive the seller agent out of the market. They are often referred as "bad-mouthing".

We propose a personalized approach that addresses unfair ratings of sellers provided by advisors but with more flexibility for buyers to weight the value of their private and public knowledge of these advisors. Our aim is to develop improved methods for modeling trustworthiness of advisors by tracking ratings provided according to their related time windows. In so doing, our approach is able to avoid the situation where advisor agents may untruthfully rate seller agents for a large number of times (known as "flooding") and deal with changes of agents' behavior. More specifically, the personalized approach for modeling the trustworthiness of an advisor agent first calculates what we refer to as the "private reputation" of the advisor, based on the buyer and advisor agents' ratings for their commonly rated seller agents. When the buyer agent is not confident in its private reputation ratings it can also

* Corresponding author. Tel.: +1 519 8884567.

E-mail addresses: j44zhang@uwaterloo.ca (J. Zhang), rcohen@uwaterloo.ca (R. Cohen).

use what we refer to as the “public reputation” of the advisor agent. This public reputation is estimated based on the advisor agent’s ratings for all seller agents in the system. The personalized approach ultimately computes a weighted combination of private and public reputations to represent the trustworthiness of the advisor. Similarly, the personalized approach for modeling the trustworthiness of a seller agent first models private reputation of the seller based on the buyer’s own ratings for the seller. If the buyer agent does not want to rely fully on its personal experience with the seller, it will consider ratings provided by advisors. It then can derive a public reputation of the seller from these ratings. Once more, a weighted combination of private and public reputations is used to determine the trustworthiness of the seller.

The essential component of the buyer’s decision making about sellers is the effectiveness of its modeling of advisor agents. We carry out experiments to demonstrate the effective value of the personalized approach in terms of adjusting advisor agents’ trustworthiness based on the percentages of unfair ratings they provided. We also show how buyers can effectively model trustworthiness of sellers, making use of advisors’ models created through the personalized approach. Our personalized model can therefore be seen as a valuable approach to use when introducing social networks in order to model the trustworthiness of sellers in electronic marketplaces.

The rest of the paper is organized as follows. In Section 2, we survey different approaches for handling unfair ratings. In Section 3, we propose a personalized approach for handling unfair ratings in an enhanced centralized trust and reputation system. Section 4 provides examples that go through each step of our approach. Experimental results are presented in Section 5. Conclusions and future work are outlined in Section 6.

2. Related work

In this section, we summarize different approaches for handling unfair ratings of sellers provided by advisors, and present our proposed categorization of those approaches.

2.1. Approaches for handling unfair ratings

Dellarocas [3] uses collaborative filtering techniques to identify the nearest neighbors of a buyer agent based on their preference similarity with the buyer. Preference similarity is measured based on the number of their similar ratings for commonly rated sellers. He then proposes the *Cluster Filtering* approach to separate ratings provided by the neighbors into two clusters. The ratings in the lower rating cluster are considered as fair ratings and the ratings in the higher rating cluster are considered as unfairly high ratings. One problem about this approach is that it does not handle unfairly low ratings.

The beta reputation system (BRS) proposed by Jøsang and Ismail [7] estimates reputation of a seller agent using a probabilistic model to propagate ratings of the seller provided by multiple advisor agents. To handle unfair feedback provided by advisors, Whitby et al. [14] extend BRS to filter out those ratings that are not in the majority amongst other ones by using the *Iterated Filtering* approach. More specifically, feedback provided by each advisor agent consists of ratings supporting both good reputation and bad reputation of a seller agent, and is represented by a beta distribution. If the cumulated reputation of the seller agent falls between the lower and upper boundaries of feedback, this feedback will be considered as fair feedback. However, the *Iterated Filtering* approach is only effective when the significant majority of ratings are fair.

Chen and Singh [2] develop a general method called GM–GC. This approach is different from filtering approaches. It explicitly

computes reputations for raters. Ratings from less reputed raters will carry less weight and have less impact on accumulated reputations of objects. To determine the reputation of a rater, the GM–GC approach first calculates quality and confidence values of each rating given by the rater to an object. It then computes the cumulated quality and confidence values of all ratings for the objects in each category or subcategory. For a system with complex categorization of objects, the computation of GM–GC will be quite time consuming.

Teacy et al. [12] propose the TRAVOS model, which is a trust and reputation model for agent-based virtual organizations. This model copes with inaccurate reputation advice by accomplishing two tasks. The first task is to estimate the accuracy of the current reputation advice based on the amount of accurate and inaccurate previous advice which is similar to that advice. The second task is to adjust reputation advice according to its accuracy. The aim of this task is to reduce the effect of inaccurate advice. However, this model assumes that seller agents act consistently, which might not be true in many cases.

Wang and Vassileva [13] propose a *Bayesian network*-based trust model in a peer-to-peer file sharing system. In this system, file providers’ capabilities are evaluated by different aspects, including download speed, file quality, and file type. A naïve Bayesian network is constructed to represent conditional dependencies between the trustworthiness of file providers and the aspects. This approach assumes that the aspects of file providers’ capabilities are conditionally independent, which is sometimes unrealistic. For instance, users may prefer high quality video and picture files, but not care much about the quality of text files.

Buchegger and Boudec [1] propose a robust reputation system for mobile ad hoc networks (RRSMAN). In RRSMAN, every node in the network maintains a reputation rating and a trust rating about every other node that it cares about. The trust rating for a node represents how likely the node will provide true advice. The reputation rating for a node represents how correctly the node participates with the node holding the rating. A modified Bayesian approach is developed to update both the reputation rating and the trust rating that a node holds for another node based on evidence collected in the past. One problem of this approach is that evidence collected by a node is weighted only according to its order of being observed. Therefore, the weights of two pieces of evidence collected one month ago and one year ago are not that different as long as they have been collected one after another. Another problem is that this approach determines the preference similarity between two nodes based only on their current reputation ratings for one other node, which is certainly insufficient.

Table 1 lists features of the approaches. In this table, “Preference” means that the approach takes into account preference similarity between buyer and advisor agents – i.e. how similar the ratings of sellers are – when it copes with unfair ratings. This feature is important because agents with different preferences may have different opinions about the seller agent’s reputation, resulting in a difference in the final rating of the seller, between the two agents. “High/Low” means that the approach is able to handle unfairly high/low ratings. The feature “Varying” indicates that the approach is able to deal with changes of seller agents’ behavior.

Table 1
Features of approaches for handling unfair ratings

Approaches	Preference	High	Low	Varying
Iterated Filtering		✓	✓	✓
TRAVOS	✓	✓	✓	
Cluster Filtering	✓	✓		✓
GM–GC		✓	✓	
Bayesian Network	✓	✓	✓	
RRSMAN	≈ ✓	✓	✓	≈ ✓

For example, BRS [14] uses a forgetting factor to assign less weight to older ratings. This feature is also important. If seller agents change their behavior, even though two ratings provided within different periods of time are different, it does not necessarily mean that one of them must be unfair. In Table 1, the mark “√” indicates that an approach has the feature. The mark “≈ √” indicates that an approach has the feature, but in a limited manner. For instance, the RRSMAN approach deals with changes of agents’ behavior by dampening advisor agents’ ratings but only according to their order of being provided.

We propose a personalized approach with the aim of having all four of these features, so that buyers can make effective selection of sellers. Our approach models the trustworthiness of advisors and critiques the advice provided by them. Advisors providing unfairly high or unfairly low ratings will have smaller trust values. The advice provided by them will then carry less weight and have less impact on buyers’ decisions. Our approach takes into account preference similarity between buyer and advisor agents by comparing their ratings for commonly rated sellers. The advisors providing more ratings in common will likely have similar preferences as the buyers, and their advice will be considered heavily. Our approach can also deal with changes of sellers’ behavior by tracking ratings of the sellers provided according to their related time windows.

2.2. Categorization

We categorize the approaches for handling unfair ratings in terms of two dimensions, a “public–private” dimension and a “global–local” dimension.

2.2.1. Public versus Private

An approach for handling unfair ratings is *private* if the buyer agent estimates the trustworthiness of an advisor agent based on only its personal experience with previous ratings provided by the advisor agent. The current rating provided by the advisor agent is likely to be fair if the advisor agent’s past ratings are also fair. For example, the TRAVOS model [12] estimates the accuracy of the advisor agent’s current rating based on the amount of fair and unfair previous ratings provided by it that are similar to its current rating. An approach for handling unfair ratings is *public* if the buyer agent estimates trustworthiness of the advisor agent based on all the ratings it has supplied for any of the seller agents in the system. A rating is likely to be reliable if it is the same as/similar to most of the other ratings for the same seller agents. For example, the *Iterated Filtering* approach [14] filters out unfair ratings that are not in the majority amongst others.

2.2.2. Global versus Local

An approach is *local* if it filters out unfair ratings based on only the ratings for the seller agent currently being evaluated as a possible partner (referred to as the current seller agent). The *Cluster Filtering* approach [3] applies a divisive clustering algorithm to separate the ratings for the current seller agent into two clusters, the lower rating cluster and the higher rating cluster. The ratings in the higher rating cluster are then considered as unfair ratings. An approach for handling unfair ratings is considered as *global* if it estimates the trustworthiness of an advisor agent based on ratings for all the seller agents that the advisor agent has rated. The GM–GC approach proposed in [2] is a *global* approach.

The categorization of the approaches for handling unfair ratings is summarized in Table 2. Note that there is no approach falling in the category of “private and local”. This is simply because there is a conflict in this category. A buyer agent typically asks advice about a seller agent from an advisor agent only when it lacks personal experience with the seller agent. An approach belonging to the “private

Table 2
Categorization of approaches for handling unfair ratings

Categories	Public	Private
Global	GM–GC	TRAVOS, RRSMAN Bayesian Network
Local	Iterated Filtering, Cluster Filtering	

and local” category will evaluate the trustworthiness of the advisor agent based only on the buyer agent’s ratings and the advisor agent’s ratings for the current seller agent. The buyer agent’s limited experience with the current seller agent is certainly not sufficient for determining the trustworthiness of the advisor agent.

We can also categorize the approaches for handling unfair ratings based on the types of the reputation systems in which they have been used. There are basically two types of reputation systems in terms of their different architectures, centralized reputation systems and distributed reputation systems [8].

In centralized reputation systems, central servers collect ratings for each seller agent from buyer agents after transactions between them have taken place. Approaches used in these systems, such as *Iterated Filtering*, *Cluster Filtering* and GM–GC, do not consider the buyer agent’s personal experience with advisor agents for their commonly rated sellers. These approaches belong to the “public” category and their determination of an advisor agent’s trustworthiness does not differ for different buyer agents.

In distributed reputation systems, there is no central location for submitting ratings or obtaining advisor agents’ ratings. A buyer agent should simply request advice about a seller agent from advisor agents. Even though some of the distributed reputation systems have distributed stores for collecting ratings, it is still costly to obtain all the ratings for a seller agent. Therefore, approaches used in these systems cannot consider all agents’ ratings for the seller agent. The approaches used in distributed reputation systems, such as TRAVOS, *Bayesian Network* and RRSMAN, handle unfair ratings by estimating the trustworthiness of an advisor agent based on each individual buyer agent’s personal experience with the advisor agent’s advice. These approaches belong to the “private” category.

The model we present in the next section takes into account both buyer agents’ private experience with advisors’ advice and the public knowledge about the advisors held by the system. Therefore, our model has the advantages of the approaches used in both centralized and distributed reputation systems (in both “public” and “private” categories). This model also offers flexibility for buyers to weight their value in the private experience and the public knowledge.

3. A personalized approach

In this section, we first describe our personalized approach for modeling the trustworthiness of advisors. The approach is used as part of a centralized reputation system. This system creates a profile for each buyer agent to record ratings for each seller it has experienced. We assume that all buyers can play the role of advisors to other buyers. We assume as well that advisors provide ratings only when a transaction occurs and these are stored with the central server. The personalized approach allows a buyer agent to estimate the reputation (referred to as private reputation) of an advisor agent based on their ratings for commonly rated seller agents.¹ When the buyer agent has limited private knowledge of

¹ We call this type of reputation private reputation because it is based on the buyer agent’s own experience with the advisor agent’s advice, and is not shared with the public. The private reputation value of the advisor agent may vary for different buyer agents.

the advisor agent, the reputation (referred to as public reputation) of the advisor agent will also be considered.² The public reputation is estimated based on all ratings for the seller agents ever rated by the advisor agent. Finally, the trustworthiness of the advisor agent will be modeled by combining the weighted private and public reputations. These weights are determined based on the estimated reliability of the private reputation. Once we have presented this framework for modeling advisors, we discuss how buyers can use this advice to model the trustworthiness of sellers, retaining an approach that combines both private and public knowledge.

3.1. Modeling trustworthiness of advisor

Our personalized approach allows a buyer agent *B* to evaluate the private reputation of an advisor agent *A* by comparing their ratings for commonly rated seller agents $\{S_1, S_2, \dots, S_m\}$. For one of the commonly rated sellers $S_i (1 \leq i \leq m \text{ and } m \geq 1)$, *A* has the rating vector R_{A,S_i} and *B* has the rating vector R_{B,S_i} . A rating for S_i from *B* and *A* is binary (“1” or “0”, for example), in which “1” means that S_i is reputable and “0” means that S_i is not reputable.³ The ratings in R_{A,S_i} and R_{B,S_i} are ordered according to the time when they are provided. The ratings are then partitioned into different elemental time windows. The length of an elemental time window may be fixed (e.g. 1 day) or adapted by the frequency of the ratings to the seller S_i , similar to the way proposed in [3]. It should also be considerably small so that there is no need to worry about the changes of sellers’ behavior within each elemental time window. We define a pair of ratings (r_{A,S_i}, r_{B,S_i}) , such that r_{A,S_i} is one of the ratings of R_{A,S_i} , r_{B,S_i} is one of the ratings of R_{B,S_i} , and r_{A,S_i} corresponds to r_{B,S_i} . The two ratings, r_{A,S_i} and r_{B,S_i} , are correspondent only if they are in the same elemental time window, the rating r_{B,S_i} is the most recent rating in its time window, and the rating r_{A,S_i} is the closest and prior to the rating r_{B,S_i} .⁴ We then count the number of such pairs for S_i , $N_{S_i}^A$. The total number of rating pairs for all commonly rated sellers, N_{all}^A will be calculated by summing up the number of rating pairs for each commonly rated seller agent as follows:

$$N_{\text{all}}^A = \sum_{i=1}^m N_{S_i}^A \quad (1)$$

The private reputation of the advisor agent is estimated by examining rating pairs for all commonly rated sellers. We define a rating pair (r_{A,S_i}, r_{B,S_i}) as a positive pair if r_{A,S_i} is the same value as r_{B,S_i} . Otherwise, the pair is a negative pair. Suppose there are N_{pos}^A number of positive pairs. The number of negative pairs will be $N_{\text{all}}^A - N_{\text{pos}}^A$. The private reputation of the advisor *A* is estimated as the probability that *A* will provide reliable ratings to *B*. Because there is only incomplete information about the advisor, the best way of estimating the probability is to use the expected value of the probability. The expected value of a continuous random variable is dependent on a probability density function, which is used to model the probability that a variable will have a certain value. Because of its flexibility and the fact that it is the conjugate prior for distributions of binary events [11], the beta family of probability density functions is commonly used to represent probability distributions

of binary events.⁵ Therefore, the private reputation of *A* can be calculated as follows:

$$\alpha = N_{\text{pos}}^A + 1, \quad \beta = N_{\text{all}}^A - N_{\text{pos}}^A + 1$$

$$R_{\text{pri}}(A) = E(\text{Pr}(A)) = \frac{\alpha}{\alpha + \beta} \quad (2)$$

where $\text{Pr}(A)$ is the probability that *A* will provide fair ratings to *B*, and $E(\text{Pr}(A))$ is the expected value of the beta distribution, which is used to define the private reputation value.

When there are not enough rating pairs, the buyer agent *B* will also consider the advisor agent *A*’s public reputation. This may happen for instance in large marketplaces where buyers and advisors may not have had experience with the same sellers. The public reputation of *A* is estimated based on its ratings and ratings from other buyers for the same sellers rated by *A*. Each time *A* provides a rating $r_{A,S}$, the rating will be judged centrally as a fair or unfair rating. We define a rating for a seller agent as a fair rating if it is consistent with the majority of ratings to the seller up to the moment when the rating is provided.⁶ As before, we consider only the ratings within a time window prior to the moment when the rating $r_{A,S}$ is provided, and we only consider the most recent rating from each advisor.

Suppose that the advisor agent *A* totally provides N_{all}^A ratings. If there are N_f^A number of fair ratings, the number of unfair ratings provided by *A* will be $N_{\text{all}}^A - N_f^A$. In the same way as estimating the private reputation, the public reputation of the advisor *A* is estimated as the probability that *A* will provide fair ratings. It can be calculated as follows:

$$\alpha' = N_f^A + 1, \quad \beta' = N_{\text{all}}^A - N_f^A + 1$$

$$R_{\text{pub}}(A) = \frac{\alpha'}{\alpha' + \beta'} \quad (3)$$

which also indicates that the greater the percentage of fair ratings advisor *A* provides, the more reputable it will be.

To estimate the trustworthiness of advisor agent *A*, we combine the private reputation and public reputation values together. The private reputation and public reputation values are assigned different weights. The weights are determined by the reliability of the estimated private reputation value.

We first determine the minimum number of rating pairs needed for *B* to be confident about the private reputation value it has of *A*. Based on the Chernoff Bound theorem [10], the minimum number of rating pairs can be determined by an acceptable level of error and a confidence measurement as follows:

$$N_{\text{min}} = -\frac{1}{2\varepsilon^2} \ln \frac{1-\gamma}{2} \quad (4)$$

where ε is the maximal level of error that can be accepted by *B*, and γ is the confidence measure. If the total number of pairs N_{all}^A is larger than or equal to N_{min} , buyer *B* will be confident about the private reputation value estimated based on its ratings and the advisor *A*’s ratings for all commonly rated sellers. Otherwise, there are not enough rating pairs, the buyer agent will not be confident about the private reputation value, and it will then also consider public reputation. The reliability of the private reputation value can be measured as follows:

$$w = \begin{cases} \frac{N_{\text{all}}^A}{N_{\text{min}}} & \text{if } N_{\text{all}}^A < N_{\text{min}} \\ 1 & \text{otherwise} \end{cases} \quad (5)$$

² We call this type of reputation public reputation because it is based on the public’s opinions about the advisor agent’s advice, and it is shared by all of the public. The public reputation value of the advisor agent is the same for every buyer agent.

³ For the purpose of simplicity, we assume ratings for sellers are binary. Possible ways of extending our approach to accept ratings in different ranges will be investigated as future work. Further discussion can be found in Section 6.

⁴ We consider ratings provided by *B* after those by *A* in the same time window, in order to incorporate into *B*’s rating anything learned from *A* during that time window, before taking an action. According to the solution proposed by Zacharia et al. [17], by keeping only the most recent ratings, we can avoid the issue of advisors “flooding” the system.

⁵ More detailed discussion of the beta function can be found in [7,8].

⁶ Determining consistency with the majority of ratings can be achieved in a variety of ways, for instance averaging all the ratings and seeing if that is close to the advisor’s rating.

The trust value of A will be calculated by combining the weighted private reputation and public reputation values as follows:

$$Tr(A) = wR_{pri}(A) + (1 - w)R_{pub}(A) \tag{6}$$

It is obvious that the buyer will rely less on the public reputation value when the private reputation value is more reliable. Note that when $w = 1$, the buyer relies only on private reputation.

3.2. Modeling trustworthiness of seller

Once we have the models of advisors, we need an effective method for the buyer agent to model the trustworthiness of a seller agent, by combining the buyer’s personal experience with the seller and reputation ratings provided by the advisors. The model of BRS [7] uses the beta probability density function to aggregate the ratings of the seller provided by the buyer and multiple advisor agents. This model, however, does not allow the buyer to weight its value in its own ratings any more or less heavily than the advisors’ ratings of the seller. We argue that buyers may rely more on their personal experience with sellers. The Bayesian network-based trust model [13] updates a Bayesian network of the seller agent’s trustworthiness based on the buyer’s direct interactions with the seller and recommendations provided by advisors that have previously interacted with the seller. This model also does not weight any differently the buyer agent’s personal experience from others’ recommendations. The TRAVOS model [12] provides a method for estimating the trustworthiness of the seller based on the buyer’s personal experience with the seller and a method for estimating the reputation of the seller by aggregating advisors’ advice. They do not provide a function for combining both of these elements. This model also assumes that seller agents act consistently; therefore, it cannot deal with changes of agents’ behavior.

Our personalized approach can also be adopted to effectively model the trustworthiness of seller agents. It allows the buyer agent to model the private reputation of a seller agent based on the buyer’s own ratings for the seller. If the buyer agent does not want to rely fully on its personal experience with the seller, it will ask for advisors’ ratings of the seller agent. It then can derive a public reputation of the seller from these ratings. The trustworthiness of the seller will be modeled by combining the weighted private and public reputation values. We formalize our approach for modeling the trustworthiness of sellers as follows.

Suppose that buyer B has the rating vector $R_{B,S}$, which contains all the ratings provided by B for the seller S . The rating of “1” will be considered as a positive rating, and “0” will be considered as a negative rating. Similarly, the ratings in $R_{B,S}$ are ordered from the most recent to the oldest according to the time when they are submitted. The ratings are then partitioned into different elemental time windows $\{T_1, T_2, \dots, T_n\}$. In this case, T_1 is the most recent (current) time window. We then count the number of positive ratings $N_{pos,i}^B$ and the number of negative ratings $N_{neg,i}^B$ in each time window T_i . The private reputation of the seller S can be estimated through the beta family of probability density functions as follows:

$$R_{pri}(S) = \frac{\sum_{i=1}^n N_{pos,i}^B \lambda^{i-1} + 1}{\sum_{i=1}^n (N_{pos,i}^B + N_{neg,i}^B) \lambda^{i-1} + 2} \tag{7}$$

where $\lambda (0 \leq \lambda \leq 1)$ is a forgetting rate. The forgetting rate is also introduced by Jøsang and Ismail [7] to deal with possible changes of the seller agent’s behavior over time because old ratings will be given less weight than more recent ones. Note that when $\lambda = 1$ there is no forgetting, and when $\lambda = 0$ only the ratings that are within the current time window T_1 will be considered.

If the buyer agent B does not have enough personal experience with the seller S , it will also consider ratings provided by other buyer agents (advisors). Suppose that advisors $\{A_1, A_2, \dots, A_k\}$ have

provided ratings for the seller S . We also partition these ratings into different elemental time windows. Suppose that the advisor A_j has provided $N_{pos,i}^{A_j}$ positive ratings and $N_{neg,i}^{A_j}$ negative ratings within the time window T_i . These ratings will be discounted based on the trustworthiness of the advisor, so that the ratings from less trustworthy advisors will carry less weight than ratings from more trustworthy ones.

Jøsang [6] provides a mapping from beliefs defined by the Dempster-Shafer theory to the beta function as follows:

$$\begin{cases} b = \frac{N_{pos,i}^{A_j}}{N_{pos,i}^{A_j} + N_{neg,i}^{A_j} + 2} \\ d = \frac{N_{neg,i}^{A_j}}{N_{pos,i}^{A_j} + N_{neg,i}^{A_j} + 2} \\ u = \frac{2}{N_{pos,i}^{A_j} + N_{neg,i}^{A_j} + 2} \end{cases} \tag{8}$$

where b , d and u represent belief, disbelief and uncertainty parameters, respectively. In our case, b represents the probability that the seller is trustworthy is true, and d represents the probability of the proposition is false. Note that $b + d + u = 1$ and $b, d, u \in [0, 1]$. As also pointed out in [7,15], beliefs and disbeliefs can be directly discounted by the trustworthiness of the advisor as follows:

$$\begin{cases} b' = Tr(A_j)b \\ d' = Tr(A_j)d \end{cases} \tag{9}$$

From Eqs. (8) and (9), we then can derive a discounting function for the amount of ratings provided by the advisor A_j as follows:

$$D_{pos,i}^{A_j} = \frac{2Tr(A_j)N_{pos,i}^{A_j}}{(1 - Tr(A_j))(N_{pos,i}^{A_j} + N_{neg,i}^{A_j}) + 2} \tag{10}$$

$$D_{neg,i}^{A_j} = \frac{2Tr(A_j)N_{neg,i}^{A_j}}{(1 - Tr(A_j))(N_{pos,i}^{A_j} + N_{neg,i}^{A_j}) + 2} \tag{11}$$

where $Tr(A_j)$ is the trustworthiness of the advisor A_j , which can be calculated by using the personalized approach as presented in the earlier section. In the same way as estimating the private reputation, the public reputation of the seller S can be calculated as follows:

$$R_{pub}(S) = \frac{\left[\sum_{j=1}^k \sum_{i=1}^n D_{pos,i}^{A_j} \lambda^{i-1} \right] + 1}{\left[\sum_{j=1}^k \sum_{i=1}^n (D_{pos,i}^{A_j} + D_{neg,i}^{A_j}) \lambda^{i-1} \right] + 2} \tag{12}$$

The ratings provided by the advisors will be also discounted by the forgetting factor λ .

Similar to the way of estimating the trustworthiness of advisors, the trustworthiness of the seller agent S is estimated by combining the weighted private and public reputation values as follows:

$$Tr(S) = wR_{pri}(S) + (1 - w)R_{pub}(S) \tag{13}$$

The weight w is determined by the reliability of the estimated private reputation value as follows:

$$w = \begin{cases} \frac{N_{all}^B}{N_{min}} & \text{if } N_{all}^B < N_{min} \\ 1 & \text{otherwise} \end{cases} \tag{14}$$

where N_{min} represents the minimum number of ratings needed for the buyer B to be confident about the private reputation value it has of S . N_{min} can be calculated by using Eq. (4). N_{all}^B is the total number of ratings provided by B for the seller.

4. Examples

To illustrate how our approach models trustworthiness of advisors and sellers, this section provides examples that go

Table 3
Ratings of sellers provided by advisors

A_j	A_x					A_y					A_z				
	T_1	T_2	T_3	T_4	T_5	T_1	T_2	T_3	T_4	T_5	T_1	T_2	T_3	T_4	T_5
S_1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0
S_2	1	1	1	1	1	0	1	0	1	1	0	0	0	0	0
S_3	1	1	1	1	1	1	0	1	0	0	0	0	0	0	0
S_4	1	1	1	1	1	1	0	0	1	0	0	0	0	0	0
S_5	1	1	1	1	1	1	1	0	0	1	0	0	0	0	0

through each step of the approach. Examples are also provided to demonstrate how trust values different buyer agents have of same advisors may vary, and to show the effectiveness of our approach even when the majority of ratings are unfair. We provide a simple example to show how to model trustworthiness of sellers after we have advisor models.

4.1. Modeling trustworthiness of advisors

In an electronic marketplace, a buyer agent B needs to make a decision on whether to interact with a seller agent S_0 , which depends on how much B trusts S_0 . To model the trustworthiness of the seller S_0 when the buyer has had no or only limited experience with S_0 , B seeks advice from three advisor agents A_x , A_y and A_z that have had experience with S_0 . The pieces of advice about S_0 from A_x , A_y and A_z are ratings representing the reputation of S_0 . Before aggregating the ratings provided by A_x , A_y and A_z , the buyer agent B needs to evaluate the reliability of those ratings, which depends on the trustworthiness of the advisors A_x , A_y and A_z . Our approach effectively models the trustworthiness of advisors based on how reliable the previous ratings provided by them are.

Consider the case where the advisors A_x , A_y and A_z each has rated only the five seller agents (S_1, S_2, S_3, S_4 , and S_5). Table 3 lists the ratings provided by A_x , A_y and A_z for the five sellers.⁷ The symbol “ T_i ” represents a sequence of time windows, in which T_1 is the most recent time window. To simplify the demonstration, we assume that each advisor agent provides at most one rating within each time window. We also assume that those are the only ratings provided by them.

As can be seen from Table 4, the buyer agent B has also provided some ratings for the five sellers. The buyer agent B might have not provided any rating for some sellers within some time window. For example, it has provided only one rating for the seller S_5 , which is in the time window T_1 . We assume that the ratings provided by B are after those provided by A_x , A_y and A_z if they are within the same time window.

We compare the ratings provided by A_x , A_y and A_z in Table 3 with the ratings provided by B in Table 4. The buyer agent B has the same number of rating pairs as each advisor agent ($N_{all}^{A_j} = 15$ and $j \in \{x, y, z\}$). However, B has different numbers of positive rating pairs with A_x , A_y and A_z , which are listed in Table 5a. Accordingly, as can be seen from Table 5a, the private reputation values of A_x , A_y and A_z are different. The private reputation value of A_x is the highest and that of A_z is the lowest. This indicates that the advisor agent A_x is most likely to provide fair ratings and have similar preferences as the buyer agent B , whereas A_z most likely will lie and have different preferences than B .

To determine the public reputation of A_x , A_y and A_z , we look at their ratings for all the sellers. According to Table 3, the total number of ratings provided by each advisor agent is the same ($N_{all}^{A_j} = 25$). We also count the number of fair ratings each advisor

⁷ Table 3 lists only the ratings provided by the advisors from which buyer B asks advice about seller S_0 . Other buyers may have also provided ratings for S_1, S_2, S_3, S_4 , and S_5 , but have not rated S_0 .

Table 4
Ratings provided by the buyer agent B

T_i	T_1	T_2	T_3	T_4	T_5
S_1	1	1	1	1	1
S_2	1	1	1	1	–
S_3	1	1	1	–	–
S_4	1	1	–	–	–
S_5	1	–	–	–	–

Table 5
Private and public reputations of advisors

A_j	A_x	A_y	A_z
(a)			
$N_{pos}^{A_j}$	15	8	0
α	16	9	1
β	1	8	16
$R_{pri}(A_j)$	0.94	0.53	0.06
(b)			
$N_f^{A_j}$	25	12	0
α'	26	13	1
β'	1	14	26
$R_{pub}(A_j)$	0.96	0.48	0.04

agent provides. A rating here is considered as a fair rating when it is consistent with the majority of ratings for the seller agent within a same time window. Consider the case where all of the five seller agents are reputable and the majority of ratings are fair. In this case, a rating of 1 provided by an advisor agent will be considered as a fair rating, whereas a rating of 0 will be considered as an unfair rating. From the advisor agents’ ratings listed in Table 3, we can see that ratings provided by the advisor agent A_x are all fair, the advisor agent A_z always lies, and some of the ratings provided by the advisor agent A_y are unfair. Table 5b lists the number of fair ratings provided by each advisor agent and the corresponding public reputation value of it. From Table 5b, it is clear that the advisor agent A_x is most likely to provide fair ratings, and the advisor A_z most likely will lie.

To combine private reputation and public reputation, the weight w should be determined. The value of w depends on the values of ϵ and γ , and the number of rating pairs $N_{all}^{A_j}$, which is the same for every advisor agent in our example. Suppose we have a fixed value, 0.8 for γ , which means that the confidence value should be no less than 0.8 in order for the buyer agent to be confident with the private reputation values of advisor agents. In this case, the more errors it can accept, the more confident it is with the private reputation values of advisor agents, which also means that the more weight it will put on the private reputation values. Table 6 lists different acceptable levels of errors, their correspondent weights of private reputation values, and different results of trust values. It clearly indicates that A_x is the most trustworthy, and A_y is more trustworthy than A_z . As a result, the buyer agent B will place more trust in the advice provided by A_x . It will consider the advice provided by A_x more heavily when aggregating the advice provided by A_x , A_y and A_z for modeling the reputation of the seller

Table 6
Trustworthiness of advisor agents

ϵ	0.1	0.15	0.2
N_{min}	115	51	29
w	0.13	0.29	0.52
$Tr(A_x)$	0.957	0.954	0.950
$Tr(A_y)$	0.487	0.495	0.506
$Tr(A_z)$	0.043	0.046	0.05

Table 7
Ratings provided by the buyer agent B'

T_i	T_1	T_2	T_3	T_4	T_5
S_1	1	1	–	–	1
S_2	1	–	–	1	–
S_3	1	1	–	–	–
S_4	1	1	–	–	–
S_5	1	–	–	–	–

Table 8
Trust values B' has of advisors

A_j	A_x	A_y	A_z
$R_{pri}(A_j)$	0.92	0.58	0.08
$R_{pub}(A_j)$	0.96	0.48	0.04
$Tr(A_j)$	0.947	0.514	0.054

Table 9
Public reputations of advisors when majority of ratings are unfair

A_j	A_x	A_y	A_z
$N_f^{A_j}$	0	13	25
α'	1	14	26
β'	26	13	1
$R_{pub}(A_j)$	0.04	0.52	0.96

agent S_0 . Our personalized approach serves the purpose of representing the trustworthiness of advisors, so that this may be taken into account, when determining how heavily to rely on their advice.

To demonstrate how the trust values different buyer agents have for the same advisors may vary, we consider another buyer agent B' , that also needs to make a decision on whether to interact with a seller agent S'_0 (S'_0 may differ from S_0). The ratings provided by B' for the five seller agents are listed in Table 7. By going through the same process as above, we can calculate the trust values the buyer agent B' has of A_x , A_y and A_z , when $\epsilon = 0.2$ and $\gamma = 0.8$. The results are presented in Table 8. Comparing Table 8 with Tables 5 and 6, we can see that the private reputations the buyer agent B' has of advisors are different from those the buyer agent B has. Although the public reputations of advisors that the buyers have are the same, the trust values that the buyers have are still different.

To show the robustness of our model, now consider a case where the majority of ratings provided by advisor agents are unfair. Adjusting our earlier example, a rating of 1 provided by an advisor agent for any seller agent will now be considered as an unfair rating, whereas a rating of 0 will be considered as a fair rating. As a result, the public reputations that the buyer B has of the advisors agents A_x , A_y and A_z will be different, which can be seen from Table 9. We model the trust values the buyer agent B has of the advisors A_x , A_y and A_z , when B 's acceptable levels of errors of private reputation values are different. Results are presented in Table 10. From this table, we can see that our approach can still correctly represent the trustworthiness of advisor agents by making adjustments to rely more heavily on the private reputations.

4.2. Modeling trustworthiness of seller S_0

In this example, we demonstrate how the buyer agent B models trustworthiness of the seller agent S_0 by using our personalized approach. We assume that the buyer B has not done any business

Table 10
Trustworthiness of advisors when majority of ratings are unfair

ϵ	0.1	0.2	0.25
N_{min}	115	29	19
w	0.13	0.52	0.79
$Tr(A_x)$	0.157	0.508	0.751
$Tr(A_y)$	0.521	0.525	0.528
$Tr(A_z)$	0.843	0.492	0.249

Table 11
Ratings of S_0 provided by A_x and A_y

T_i	T_1	T_2	T_3	T_4	T_5
A_x	0	0	0	1	1
A_y	1	1	1	1	1

Table 12
Amount of Ratings of S_0 provided by A_x and A_y

T_i	T_1	T_2	T_3	T_4	T_5
$N_{pos,i}^{A_x}$	0	0	0	1	1
$N_{neg,i}^{A_x}$	1	1	1	0	0
$N_{pos,i}^{A_y}$	1	1	1	1	1
$N_{neg,i}^{A_y}$	0	0	0	0	0

with the seller S_0 . Therefore, the private reputation of S_0 can be calculated according to Eq. (7) as follows:

$$R_{pri}(S_0) = \frac{0 + 1}{(0 + 0) + 2} = 0.5$$

The buyer B then asks advice from advisors A_x , A_y and A_z . Results from the earlier examples show that the trust values that B has of advisors A_x , A_y and A_z are 0.95, 0.506 and 0.05, respectively, when we set ϵ to be 0.2. Because advisor A_z has a very low trust value, we assume that the buyer B will consider advice from only the advisors A_x and A_y .⁸

The ratings of the seller S_0 provided by the advisors A_x and A_y are listed in Table 11. We assume that the seller S_0 is trustworthy in almost half the time. We first count the amount of positive and negative ratings provided by the advisors A_x and A_y within each time window, as listed in Table 12. We then discount the amount of ratings provided by them, using Eqs. (10) and (11). The discounted amount of ratings is listed in Table 13.

In this example, we set λ to be 0.9, which means that the buyer B does not have much forgetting. According to Eq. (12), the public reputation of the seller can be calculated as follows:

$$R_{pub}(S_0) = \frac{\sum_{i=4}^5 0.927 * 0.9^{i-1} + \sum_{i=1}^5 0.406 * 0.9^{i-1} + 1}{\sum_{i=1}^5 0.927 * 0.9^{i-1} + \sum_{i=1}^5 0.406 * 0.9^{i-1} + 2} = 0.529$$

Because the buyer B has not done business with the seller before, the weight w' of the private reputation of the seller is 0. The trustworthiness of the seller S_0 can then be calculated by using Eq. (13) as follows:

$$Tr(S_0) = 0 * 0.5 + (1 - 0) * 0.529 = 0.529$$

We calculate the public reputation of the seller by taking into account the trustworthiness of advisors. From the result of $Tr(S_0)$,

⁸ What is required is then an approach for limiting the number of advisors that are consulted. For simplicity in this example, we assume some kind of threshold is used and trustworthiness of advisors must be greater than 0.05 at least. By doing so, we can cope with the situation where a buyer may falsely improve its trustworthiness by creating multiple fake identities [16].

Table 13

Discounted amount of ratings of S_0 provided by A_x and A_y

T_i	T_1	T_2	T_3	T_4	T_5
$D_{pos,i}^{A_x}$	0	0	0	0.927	0.927
$D_{neg,i}^{A_x}$	0.927	0.927	0.927	0	0
$D_{pos,i}^{A_y}$	0.406	0.406	0.406	0.406	0.406
$D_{neg,i}^{A_y}$	0	0	0	0	0

we can see that the buyer relies on the advice provided by A_x more heavily, and A_y 's advice has less impact on the result. We compare it with the way of not considering the trustworthiness of advisors. The public reputation of the seller will be calculated as follows:

$$R'_{pub}(S_0) = \frac{\sum_{i=4}^5 1 * 0.9^{i-1} + \sum_{i=1}^3 1 * 0.9^{i-1} + 1}{\sum_{i=1}^5 1 * 0.9^{i-1} + \sum_{i=1}^5 1 * 0.9^{i-1} + 2} = 0.636$$

The trustworthiness of the seller S_0 can then be calculated by as follows:

$$Tr'(S_0) = 0 * 0.5 + (1 - 0) * 0.636 = 0.636$$

From the results of $Tr(S_0)$ and $Tr^*(S_0)$, we can see that the trust value of the seller, calculated through our formula, is closer to the actual trustworthiness of the seller. It indicates that our formulation results in better estimation for the trustworthiness of the seller.

5. Experimental results

Our approach models the trustworthiness of advisors according to the reliability of the ratings provided by them. To demonstrate the effectiveness of the approach, we carry out experiments involving advisors that provide different percentages of unfair ratings. The expectation is that trustworthy advisors will be less likely to provide unfair ratings, and vice versa. We also examine how large numbers of dishonest advisors (i.e. advisors that provide unfair ratings) will affect the estimation of advisors' trustworthiness. Results indicate that our approach is still effective by making adjustments to rely more heavily on private reputations of advisors, in this case. We conduct further experiments to test the scalability of our approach. Results show that trustworthiness of advisors remains nearly the same for different populations of sellers. We also demonstrate how buyers can effectively model trustworthiness of sellers using the personalized approach, making use of advisors' models.

The first experiment involves 100 sellers, 3 buyers, and one advisor. The 3 buyers, B_1 , B_2 and B_3 , rate 10, 40 and 70 randomly selected sellers, respectively. The advisor totally rates 40 randomly selected sellers.⁹ We examine how the trust values the buyers have of the advisor change when different percentages (from 0% to 100%) of its ratings are unfair. As illustrated in Fig. 1, the trust values the buyers have of the advisor decrease when more percentages of the advisor's ratings are unfair. From this figure, we can also see that our approach is still effective when the buyer B_1 does not have much experience with sellers, in the sense that B_1 can still reduce the reputation of the advisor when it provides more unfair ratings.

The second experiment involves 100 sellers, 80 advisors, and one buyer. The buyer and each advisor rate 80 of the randomly selected sellers. We model the trust value the buyer has of one of the advisors, A. The trustworthiness of the advisor will be modeled as the combination of its private and public reputations (referred to as the CR approach) and as only its public reputation (referred to as the PR approach), respectively. The advisor A will provide different percentages (from 10% to 100%) of unfair ratings. Fig. 2

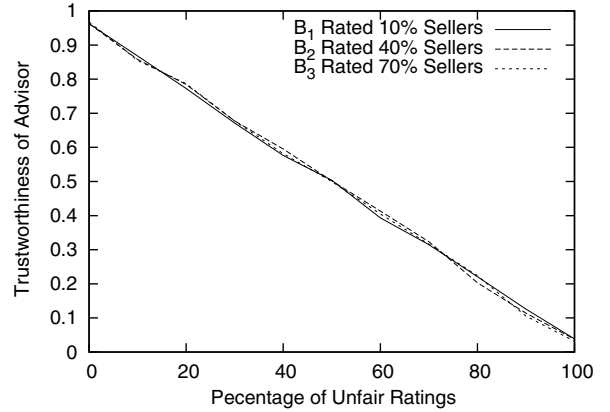


Fig. 1. Trustworthiness of advisor.




Fig. 2. Trustworthiness of A when majority of advisors are honest.

illustrates the trustworthiness of A when 24 (30% of all) advisors are dishonest. Those dishonest advisors provide the same percentage of unfair ratings as the advisor A does. Results indicate that the trustworthiness of A modeled by using the CR and PR approaches decreases when more percentages of ratings provided by A are unfair. Therefore, these two approaches are not affected when only a small number of advisors are dishonest. Fig. 3 represents the trustworthiness of A when 48 (60% of all) advisors are dishonest. In this figure, the trustworthiness of A modeled by using the CR approach still decreases when more percentages of ratings provided by A are




Fig. 3. Comparison of the CR and PR approaches.

⁹ Note that we simplify the experiments by limiting each buyer or advisor to provide at most one rating for each seller.

unfair, which indicates that our approach is still effective when the majority of advisors provide large numbers of unfair ratings. In contrast, the trustworthiness modeled by using the PR approach increases when more than 60% of ratings provided by the dishonest advisors are unfair, which indicates that the PR approach is only effective when the majority of ratings are fair. The statistical significance of the results is also confirmed in the figure by the fact that the intervals (corresponding to ± 1 standard deviation) do not overlap when more than 60% of ratings provided by the dishonest advisors are unfair.

The effectiveness of our approach is demonstrated by the above experiments with the fixed population of (100) sellers. It is useful to examine whether our approach will still be useful when there are a large number of sellers. The number of sellers affects the number of commonly rated sellers, and may then affect the calculation of private reputation for advisors. More specifically, in the environment where there are many sellers, there may be a smaller percentage of those sellers that have been commonly rated by buyers and advisors. In this case, buyers may have less private knowledge about advisors. We use a simulation to demonstrate that our approach can still effectively model trustworthiness of advisors. In this simulation, we have different populations of sellers spanning from 100 to 500 in increments of 50. A buyer models trustworthiness of an advisor. 50% of the ratings provided by the advisor are unfair in this experiment. The results are shown in Fig. 4. The x-axis represents the populations of sellers, and the y-axis represents the trustworthiness of the advisor. The solid line is the average trust value of the advisor. As can be seen from Fig. 4, the trustworthiness of the advisor remains nearly the same when the population of sellers changes, which indicates that our approach is scalable.

After demonstrating the effectiveness of our approach in modeling trustworthiness of advisors, we carry out a further experiment to examine how buyers can make use of our method for modeling advisors in order to effectively model the trustworthiness of sellers. This experiment also involves 100 sellers, 80 advisors, and one buyer. Similarly, the buyer and each advisor rate 80 of the randomly selected sellers. Every 10% of the sellers acts dishonestly with different probabilities (from 0 to 0.9). The buyer models the trustworthiness of sellers based on the advisors' ratings of sellers. In order to determine which advisors the buyer should ask advice from, the buyer first models trustworthiness of advisors, and then selects a list of trustworthy advisors from which it can ask advice about sellers. Once this list is determined, the ratings of each of the advisors in the list need to be combined to determine the trustworthiness of the sellers. For this experiment, we assume that only the 10 most trustworthy advisors are consulted.¹⁰

Similar to the second experiment, the trustworthiness of each advisor will be modeled based on either the CR approach or the PR approach. Fig. 5 illustrates the trustworthiness of different sellers when 30% of advisors are dishonest. Results indicate that the trustworthiness of sellers, when using the CR and PR approaches to model trustworthiness of advisors, decreases when they act dishonestly with higher probabilities. Therefore, these two approaches are both effective when only a small number of advisors are dishonest. Fig. 6 represents the trustworthiness of sellers when 60% of advisors are dishonest. In this figure, the trustworthiness of sellers, when using the CR approach to model trustworthiness of

¹⁰ Note that other methods may be used to determine the list of trustworthy advisors to consult (for example, using a threshold and retaining only advisors with trustworthiness beyond that threshold). Also note that larger list will increase computation, and may decrease the accuracy for predicting seller agents' trustworthiness from advice provided by advisors. A smaller list may increase the accuracy, but will have higher chance that none of the advisors has rated some sellers. The detailed study of the proper number of advisors needed to be consulted can be found in [5].

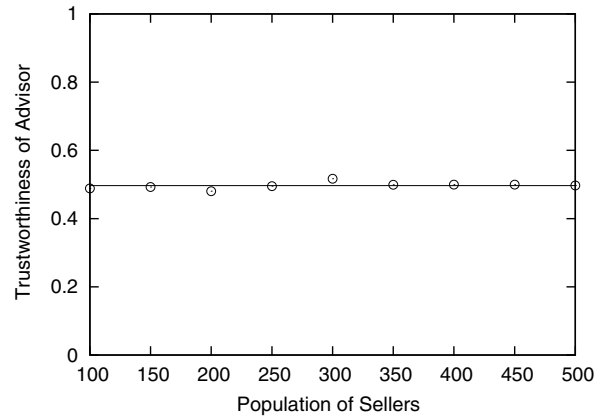


Fig. 4. Scalability of our approach.




Fig. 5. Trustworthiness of sellers when majority of advisors are honest.




Fig. 6. Comparison of the CR and PR approaches.

advisors, still decreases when the sellers act dishonestly in higher probabilities, which indicates that our approach is still effective when the majority of advisors provide large numbers of unfair ratings. In contrast, the trustworthiness of sellers, when using the PR approach to model trustworthiness of advisors, increases when the sellers act dishonestly in higher probabilities. This indicates that the PR approach is only effective when the majority of ratings are fair. This figure also shows that the intervals do not overlap,

which confirms the statistical significance of our results. All in all, if taking our model and using it as a basis for evaluating sellers, more accurate decisions about trustworthiness of sellers can be made than using other methods for modeling advisors.

Note that we do not provide experiments to demonstrate how the trustworthiness of sellers will change when the population of agents changes. As demonstrated by the experiments, our personalized approach can effectively model the trustworthiness of advisors. By also using the personalized approach for modeling the trustworthiness of sellers, buyers can always effectively adjust ratings provided by advisors based on the trustworthiness of the advisors. Therefore, our approach should also be able to scale well when modeling the trustworthiness of sellers.

6. Conclusions and future work

In this paper, we first survey different approaches for handling unfair ratings, and categorize them according to various features. Approaches for handling unfair ratings should be able to take into account the preference similarity between buyer agents and advisor agents. They should be able to handle both unfairly high and low ratings. They should also be able to deal with changes of agents' behavior over time. We then also categorize these approaches in terms of two dimensions, a "public-private" dimension and a "global-local" dimension, and categorize them based on the types of reputation systems in which they have been used. Approaches used in centralized reputation systems belong to the "public" category and do not consider buyer agents' personal experience with advisor agents' advice (ratings), whereas approaches used in distributed reputation systems belong to the "private" category and cannot consider all ratings for seller agents. This categorization of the different approaches provides a valuable perspective on the key challenges faced in designing an effective reputation system that makes use of advice from other agents, but takes care to consider the trustworthiness of those ratings.

Based on the study of these approaches, we propose a personalized approach for effectively handling unfair ratings in centralized reputation systems. The personalized approach has all of the desirable features that we outlined. It also has the advantages of both approaches used in centralized reputation systems and approaches used in distributed reputation systems. It allows a buyer agent to estimate the private reputation of an advisor agent based on their ratings for commonly rated seller agents. When the buyer agent is not confident with the private reputation value, it can also use the public reputation of the advisor agent. The public reputation of the advisor agent is evaluated based on all ratings for the seller agents rated by the advisor agent. Similarly, we adopt the personalized approach to model the trustworthiness of seller agents by combining the weighted private and public reputation values of the sellers. Experimental results demonstrate the effectiveness of the personalized approach in terms of adjusting agents' trustworthiness based on the percentages of unfair ratings they provided. Trustworthiness of advisor agents will be decreased more/less if advisor agents provide more/fewer unfair ratings. Our approach can effectively model the trustworthiness of advisors even when buyer agents do not have much experience with seller agents. Furthermore, our approach is still effective when the majority of advisor agents provide large numbers of unfair ratings, by adjusting to rely more heavily on private reputations of advisor agents. In addition we show that our approach is scalable in terms of different populations of involved sellers. We also demonstrate the value of our method for modeling advisors in order to effectively model the trustworthiness of sellers.

For future work, we will also carry out further experiments to continue to compare our model with competing approaches, such

as the BRS model and the TRAVOS model. The performance of the approaches could be evaluated, for instance, based on average estimation error, which is the average difference between seller agents' actual trust values and estimated trust values. In fact, we would like to see our personalized approach provide for an attractive environment in which to conduct business: allowing agents to represent either consumers or businesses, as they operate with improved procedures for interpreting the information they receive about sellers from other buyer agents.

Another avenue for future work is to make adjustments to the current model, to broaden its applicability. In the current framework, we allow buyers to be advisors and do not consider explicitly how to create incentives for the advisors to not only report their ratings but to do so truthfully. We do have methods in place for modeling trustworthiness, but it would be even more beneficial if advisors were motivated to be honest. We have in fact conducted some preliminary research into creating incentives for honesty, based on rewards that sellers offer to buyers that are well accepted advisors in the social network. This work is reported in [18].

Another possible extension is to move beyond binary ratings for seller agents to accept ratings in different ranges. In this case, we could begin with a modest set of possible values, each with a qualitative interpretation (e.g. very reputable, neutral, not reputable, etc.) as in [2]. The Dirichlet family of probability density functions [4], which is the multivariate generalization of the beta family, can be used to represent probability distributions of discrete values. Another possible extension is to allow advisors and buyer agents to represent the reputation of a seller agent not as a single rating but as a rating of different dimensions of trustworthiness. We could, for example examine different aspects (e.g. delivery time, quality and prices) of sellers' products similar as used by Wang and Vassileva [13], but take into account relationships among those aspects by using for example, a quality of service ontology used by Maximilien and Singh [9].

At the moment, we have set aside the question of how best to determine the appropriate size of the time window to be used in the evaluation of the trustworthiness of advisors. The time windows proposed in [3], for instance, are determined based on the frequency of the ratings to a given seller, so that if the market carries many ratings to this seller the time windows are quite small. This suggests that some methods for gauging the level of activity in the market could be applied to the proposed size of time window. Experiments could be conducted to examine the relative benefits of different sizes of time windows, in effectively capturing changes of sellers' behavior.

Another potential future work is to distinguish ratings for the current seller agent from ratings for other seller agents. As stated earlier in the related work section, there is no approach belonging to the "private and local" category because buyer agents' limited experience with the current seller agent is insufficient to estimate trustworthiness of advisor agents. However, we believe that ratings for the current seller agent should influence buyer agents' decisions more heavily, and therefore should gain more weight when estimating trustworthiness of advisor agents.

A final area deserving further study is how best to determine which advisors to consult, when modeling the trustworthiness of sellers. Challenging problems in this area include how to benefit from the greater information source when the number of advisors is large, but temper this by the need to address the greater chance for unreliability when the pool of advisors is not small.

Acknowledgement

Financial assistance was received from NSERC (Natural Science and Engineering Research Council of Canada).

References

- [1] S. Buchegger, J.-Y.L. Boudec, A robust reputation system for mobile ad-hoc networks, Technical Report, IC/2003/50, EPFL-IC-LCA, 2003.
- [2] M. Chen, J.P. Singh, Computing and using reputations for internet ratings, in: Proceedings of the Third ACM Conference on Electronic Commerce, 2001, pp. 154–162.
- [3] C. Dellarocas, Immunizing online reputation reporting systems against unfair ratings and discriminatory behavior, in: Proceedings of the Second ACM Conference on Electronic Commerce, 2000, pp. 150–157.
- [4] A. Gelman, J.B. Carlin, H.S. Stern, D.B. Rubin, Bayesian Data Analysis, Chapman & Hall/CRC, Boca Raton, 2004.
- [5] J. Herlocker, J.A. Konstan, J. Riedl, An empirical analysis of design choices in neighborhood-based collaborative filtering algorithms, *Information Retrieval* 5 (4) (2002) 287–310.
- [6] A. Jøsang, A logic for uncertain probabilities, *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* 9 (3) (2001) 279–311.
- [7] A. Jøsang, R. Ismail, The Beta Reputation System, in: Proceedings of the 15th Bled Electronic Commerce Conference, 2002, pp. 324–337.
- [8] A. Jøsang, R. Ismail, C. Boyd, A survey of trust and reputation systems for online service provision, *Decision Support Systems* 43 (2) (2005) 618–644.
- [9] E.M. Maximilien, M.P. Singh, Agent-based trust model using multiple qualities, in: Proceedings of the Fourth International Autonomous Agents and Multiagent Systems (AAMAS'05), 2005, pp. 519–526.
- [10] L. Mui, M. Mohtashemi, A. Halberstadt, A computational model of trust and reputation, in: Proceedings of the 35th Hawaii International Conference on System Science (HICSS'02), vol. 7, 2002, pp. 2431–2439.
- [11] S. Russell, P. Norvig, *Artificial Intelligence: A Modern Approach*, Second ed., Prentice Hall, Englewood Cliffs, New Jersey, 2002.
- [12] W.T.L. Teacy, J. Patel, N.R. Jennings, M. Luck, Coping with inaccurate reputation sources: experimental analysis of a probabilistic trust model, in: Proceedings of the Fourth International Autonomous Agents and Multi Agent Systems (AAMAS'05), 2005, pp. 997–1004.
- [13] Y. Wang, J. Vassileva, Bayesian network-based trust model in peer-to-peer networks, in: Proceedings of the Sixth International Workshop on Trust, Privacy, Deception and Fraud in Agent Systems, 2003, pp. 23–34.
- [14] A. Whitby, A. Jøsang, J. Indulska, Filtering out unfair ratings in bayesian reputation systems, *The Icfain Journal of Management Research* 4 (2) (2005) 48–64.
- [15] B. Yu, M. Singh, Detecting deception in reputation management, in: Proceedings of International Autonomous Agents and Multiagent Systems (AAMAS'03), 2003, pp. 73–80.
- [16] G. Zacharia, P. Maes, Trust management through reputation mechanisms, *Journal of Applied Artificial Intelligence* 14 (9) (2000) 881–907.
- [17] G. Zacharia, A. Moukas, P. Maes, Collaborative reputation mechanisms in electronic marketplaces, in: Proceedings of the 32nd Hawaii International Conference on System Sciences (HICSS'99), vol. 8, 1999, pp. 8026–8032.
- [18] J. Zhang, R. Cohen, Design of a mechanism for promoting honesty in e-marketplaces, in: Proceedings of the 22nd National Conference on Artificial Intelligence (AAAI'07), 2007, pp. 1495–1500.