A Reputation Revision Mechanism to Mitigate the Negative Effects of Misreported Ratings

Siyuan Liu[1]   Chunyan Miao[1,2]   Yuan Liu[1]   Hui Fang[1]   Han Yu[1]   Jie Zhang[2]   Cyril Leung[1,3]

[1]Joint NTU-UBC Research Centre of Excellence in Active Living for the Elderly
[2]School of Computer Engineering
Nanyang Technological University, Singapore
[3]Electrical and Computer Engineering, The University of British Columbia, Canada
{syliu, ascymiao, liu.yuan, hfang1, han.yu, zhangj, cleung}@ntu.edu.sg

ABSTRACT

Reputation systems aggregate the ratings provided by buyers to gauge the reliability of sellers in e-marketplaces. The evaluation accuracy of seller reputation significantly impacts the sellers' future utility. The existence of unfair ratings is well-recognized to negatively affect the accuracy of reputation evaluation. Most of the existing approaches dealing with unfair ratings focus on filtering/discounting/aligning the possible unfair ratings caused by malicious attacks or subjective difference. However, these approaches are not effective against unfair ratings in the form of misreporting (e.g., a well-behaving buyer misjudged a seller and provided a negative rating to a transaction which deserves a positive one, and the buyer is willing to revert the misreported negative rating). In this case, how should the buyer undo the damage caused by such misreported ratings and help the seller recover utility loss? In this paper, we propose a reputation revision mechanism to mitigate the negative effects of the misreported ratings. The proposed mechanism temporarily inflates the reputation of the misjudged seller for a period of time, which allows the seller to recover his utility loss caused by the misreported ratings. Extensive realistic simulation based experiments demonstrate the necessity and effectiveness of the proposed mechanism.

CCS Concepts

•Information systems → Information systems applications; Collaborative and social computing systems and tools; Reputation systems;

Keywords

Misreport, Reputation, Mechanism, Revenue

1. INTRODUCTION

Reputation systems [25] are proposed to assist a buyer in selecting trustworthy transaction partners in e-marketplaces.

By making an impact on the future expected utility of a seller, reputation systems can be viewed as a sanctioning mechanism for a community to self-police. Thus, the accuracy of reputation evaluation is important for the well-being of both buyers and sellers. One of the challenges faced by today's reputation systems is that of unfair ratings [24]. Computational approaches have been proposed to address the unfair ratings caused by malicious attacks or consistent subjective difference. However, these existing approaches have not considered the case in which a well-behaving buyer misjudges a seller and provides a negative rating to a transaction that deserves a positive rating, and the buyer is willing to revert the misreported negative rating. This situation is referred to as *misreporting*, which can be caused by unintentional factors such as miscommunications illustrated by the following example:

In an e-marketplace (e.g., eBay [6]), Alice provided a negative rating to a seller, Bob, because it appeared to her that she did not receive the ordered item on time. Bob's reputation dropped accordingly, and other buyers adapted their decisions in view of this change in Bob's reputation. Several days later, Alice found out she had made a mistake in providing rating to Bob as her mother had actually signed for the item (which arrived on time) on behalf of her, but forgot to pass her the item.

Compared with the unfair ratings caused by malicious attacks or subjective difference, the misreported ratings have the following features: 1) the misreported ratings are occasionally provided by well-behaving buyers, which are different from the intentional unfair ratings caused by malicious attacks; and 2) the buyer is willing to revert the negative effects of the misreported ratings, which is different from the situation that the buyers provide intentional unfair ratings or unintentional subjective ratings.

Although no specific data is available on how widespread the problem of misreporting is, it is frequently observed that the sellers complained that buyers provided wrong ratings which caused utility loss as a consequence [1, 2, 4, 5]. It is apparently significant enough to prompt major e-marketplace operators to implement mechanisms to address this issue. Current approaches are often based on the intuition drawn from the trusting behaviors arising from face-to-face scenarios. For example, in the most popular Chinese e-marketplace – Taobao [7], buyers are allowed to provide additional comments to transactions they have already rated. However, the

dissemination of the additional remedial comments is not efficient as buyers who have viewed the misreported ratings might not return to read the additional comments. Furthermore, the seller who received the misreported ratings still suffered utility loss before the buyer realized that he has misjudged the seller.

The sophisticated approaches to addressing unfair ratings involve statistically or probabilistically filtering or discounting the possible unfair ratings [10, 13, 14, 20, 21, 22, 26], or aligning the subjective ratings through buyer behavior modelling or learning [17, 19, 23]. However, these approaches focus on addressing the unfair ratings caused by malicious attacks or subjective difference, and thus are not applicable to address the problem of misreporting. For example, a seller received 100 ratings, of which 90 are positive and 10 are negative. The coming of 1 misreported negative rating will be failed to be addressed by the existing approaches as this misreported rating does not significantly change the statistical/probalistic pattern of the ratings received by the seller, or it does not violate the consistency of the subjective behavior of the buyer in providing ratings. As a consequence, the misreported rating will make the seller reputation inaccurately evaluated, causing the seller to suffer utility loss.

In this paper, we propose a reputation revision mechanism, *RepRev*, to recover the sellers' utility loss caused by misreported ratings. RepRev temporarily inflates the misjudged seller's reputation for a period of time. The reputation inflation value and the inflation period are determined by the number of the misreported ratings and the estimated utility loss during the time between the provision and discovery of the misreports. Therefore, RepRev contributes to restoring the seller's reputation and compensating the seller's utility loss. We conduct extensive simulations based on realistic settings to demonstrate that the misreported ratings negatively impact the sellers' utility, and RepRev can mitigate the adverse effects through compensating the sellers' utility loss.

2. RELATED WORK

In this section, we briefly review the existing approaches to addressing the problem of unfair ratings and the available mechanisms to prevent and revise misreported ratings.

The existing approaches to addressing the problem of unfair ratings are mostly proposed in two directions: 1) probabilistically or statistically filtering or discounting unfair ratings; and 2) modelling or learning consistent buyer behaviors to align their ratings when evaluating seller reputation.

2.1 Probabilistic and Statistical Approaches

From the aspect of filtering unfair ratings, the iterated filtering approach [22] was proposed to filter unfair ratings for Beta Reputation System (BRS) [12] if a buyer's ratings are outside the q or $1 - q$ quantile of the majority buyers' ratings. In the entropy based approach [21], a buyer's ratings are filtered if these ratings deviate from other buyers' ratings, where the deviation is measured based on an entropy based metric. Clustering techniques are also used to filter unfair ratings. In [8], the buyers' ratings are separated into two clusters, and the ratings in the cluster containing lower ratings are considered as unfair ratings. In [10], a two-layered clustering approach was proposed. More specifically, in the first layer, the unfair ratings that are probably caused by malicious attacks are filtered out. In the second layer, the

unfair ratings that are possibly caused by subjective difference are aligned.

From the aspect of discounting unfair ratings, TRAVOS [20] was proposed to discount a buyer's ratings according to the historical accuracy of the ratings provided by the buyer. The personalized approach [26] measures the reliability of a buyer in providing ratings and discount the buyer's ratings according to the measurement. The fuzzy logic based approach [14] measures the reliability of a buyer's ratings by considering the buyer's expertise, the similarity between the buyer and other buyers, and the time when the ratings are provided. The Dempster-Shafer theory was also used to discount a buyer's ratings [13].

The statistical or probabilistic approaches rely on constructing a pattern of the ratings provided by a buyer to filter or discount possible unfair ratings. Therefore, these approaches may fail to address misreported ratings as these occasionally happened ratings may not change the pattern.

2.2 Learn Buyers' Behavior

In [17] and [19], a buyer's ratings are aligned through learning the buyer's behaviors using a Bayesian network. The Bayesian approaches learn a buyer's behavior pattern and infer this pattern back to ratings based on the observed similarities between groups of buyers. In [23], the authors proposed a reinforcement learning based reputation model which adjusts the relative importance given to the ratings from each buyer based on the actual gain or loss derived from the transactions following their recommendations.

The approaches learning buyer behaviors depend on consistent buyer behavior pattern to align a buyer's ratings. It is difficult to apply such approaches to address the problem of misreporting as these misreported ratings may not violate the buyer's consistent behaviors. For example, the misreported negative ratings from an optimistic buyer may be failed to be aligned as positive.

2.3 Prevent and Revise Misreports

In practice, some mechanisms have been adopted by current e-marketplaces to prevent the happening of misreported ratings or to revise misreported ratings once they happened. For instance, the mechanism adopted by Taobao[1] is to provide a 30-minutes buffer between the time when a rating is reported by a buyer and the time when the rating is made public. During the buffer time, the buyer still can make changes to his rating. The seller can use this time to communicate with the buyer to eliminate misunderstandings if there are any. Nevertheless, 30 minutes is a short time to effectively resolve misunderstandings such as in the example situation illustrated in the previous section.

In eBay, a seller has 5 chances to request feedback revision among 1000 feedbacks in consecutive 12 months [3]. After going through a complex procedure, the seller may get the misreports removed, but he still suffers utility loss during the period.

To the best of our knowledge, the proposed mechanism is the first to provide the capability of compensating sellers' utility loss due to misreported ratings. Our work differs from the existing approaches in the sense that we are dealing with

---

[1]Taobao was launched in 2003 by the Alibaba Group, Inc. and has now become the most popular e-marketplace in China. By January 2012, Taobao had 180 million registered users, including 2 million sellers.

the negative effects of the misreported ratings, not only on sellers' reputation, but also on sellers' utility.

3. THE PROPOSED MECHANISM

In this paper, we focus on compensating a seller's utility loss caused by misreported negative ratings, i.e., the negative ratings that are wrongly attributed to transactions that deserve positive ratings. The factors that motivate us to focus on misreported negative ratings are as follows. Firstly, negative ratings make a deep impression. Reputation is often difficult to build but easy to destroy. In this sense, a negative rating carries a heavier weight than a positive one. Secondly, compared to positive ratings, negative ratings are more likely to be widely disseminated. As a Chinese saying aptly puts it – "good news tends to stay indoors while bad news often goes far away". Thirdly, it is seldom observed that a buyer wants to revert his provided positive rating. Therefore, in this work, we focus on compensating a seller's utility loss when misreported negative ratings happen. Continuing the example mentioned in Section 1, a typical scenario to initiate the procedure of seller utility compensation can be as follows:

Alice has found out that she misjudged Bob and provided a wrong negative rating. She wants to revert the effect of this negative rating. Therefore, she reports this to the system, and the system initiates the procedure of compensating Bob after Alice and Bob reach a mutual understanding.

Assuming that a seller's future expected utility monotonically increases with his reputation, the proposed RepRev mechanism compensates the seller's potential utility loss by tentatively increasing the seller's reputation for a period of time. Suppose a seller s receives γ_s misreported ratings at time t_0. Then, s's reputation may drop by as much as δ_s. Let $R_s(t)$ denote s's reputation at time t if the misreports did not occur, and $R'_s(t)$ denote s's reputation at time t with the misreports. The exact value of δ_s depends on a given reputation system. For example, under the Beta Reputation System (BRS) [12] (in BRS, a seller's reputation is evaluated as the expected value of a positive outcome happening in the future following a Beta distribution, whose parameters are the numbers of the positive and negative ratings that the seller received in the past), δ_s can be calculated as:

$$\delta_s = R_s(t_0) - R'_s(t_0) = \frac{(\alpha_s(t_0) + \gamma_s) + 1}{(\alpha_s(t_0) + \gamma_s) + (\beta_s(t_0) - \gamma_s) + 2}$$
$$- \frac{\alpha_s(t_0) + 1}{\alpha_s(t_0) + \beta_s(t_0) + 2} \quad (1)$$
$$= \frac{\gamma_s}{\alpha_s(t_0) + \beta_s(t_0) + 2},$$

where $\alpha_s(t_0)$ and $\beta_s(t_0)$ denote the total number of positive and negative ratings (including γ_s misreported negative ratings) received by s until t_0, respectively.

As another example, in the reputation system adopted in Taobao (which is similar to the one used by eBay), s receives a score of 1 for a positive rating, and -1 for a negative rating. Considering s's reputation as his accumulated rating score, δ_s can be calculated as:

$$\delta_s = R_s(t_0) - R'_s(t_0) = (\alpha_s(t_0) + \gamma_s) - (\beta_s(t_0) - \gamma_s)$$
$$- (\alpha_s(t_0) - \beta_s(t_0)) = 2\gamma_s. \quad (2)$$

To compensate s's utility loss, RepRev first restores s's reputation from $R'_s(t)$ to $R_s(t)$, then it temporarily inflates s's reputation by δ_s for a period of time, as illustrated in Figure 1 (this figure servers as example only), where the x-axis represents time, $U(R_s(t))$ is the utility function with respect to s's reputation at t. In practice, this function can be determined through statistical analysis as the past transaction information in many large scale e-marketplaces is recorded. In Figure 1, t_0 is the time when the misreports happen, t_1 is the time when the misreports are found and compensation procedure is initiated, and t_2 is the time when the compensation completes. There are three curves shown in the figure:

- $U(R_s(t))$: the utility function with respect to s's reputation when there are no misreports.

- $U(R'_s(t))$: the utility function with respect to s's reputation when there are misreports.

- $U(R_s(t) + \delta_s)$: the utility function with respect to s's inflated reputation when s's reputation without misreports is inflated with δ_s on the basis of $R_s(t)$.

The solid parts of the utility function curves represent the actual utility of seller s, and the dash parts represent the supposed trends of the above three utility functions. For example, from the beginning to t_0, s's utility function is $U(R_s(t))$. From t_0 to t_1, s's utility function is $U(R'_s(t))$. From t_1 to t_2, s's utility function is $U(R_s(t) + \delta_s)$. After t_2, s's utility function turns back to $U(R_s(t))$.


Figure 1: The demonstration of the proposed mechanism

More specifically, suppose the buyer becomes aware of the misreported ratings at time t_1. It can be seen that s may have suffered an expected utility loss Δ_1 from time t_0 to t_1. Assuming that $U(R_s(t))$ monotonically increases with $R_s(t)$, Δ_1 can be calculated as:

$$\Delta_1 = \int_{t_0}^{t_1} U(R_s(t))dt - \int_{t_0}^{t_1} U(R'_s(t))dt. \quad (3)$$

Then RepRev begins to compensate s's expected utility loss by temporarily inflating s's reputation to $R_s(t) + \delta_s$. The inflation will end at t_2. The expected compensated utility Δ_2 can be calculated as:

$$\Delta_2 = \int_{t_1}^{t_2} U(R_s(t) + \delta_s)dt - \int_{t_1}^{t_2} U(R_s(t))dt. \quad (4)$$

Ideally, the expected compensated utility should be equal to the expected utility loss:

$$\Delta_1 = \Delta_2. \tag{5}$$

Substituting Eqs. (3) and (4) into Eq. (5), t_2 can be theoretically approximated as:

$$t_2 = \frac{\int_{t_0}^{t_1} U(R_s(t))dt - \int_{t_0}^{t_1} U(R'_s(t))dt}{U(R_s(t_1) + \delta_s) - U(R_s(t_1))} + t_1. \tag{6}$$

Here, we approximate $U(R_s(t))$ with $U(R_s(t_1))$ where $t \in [t_1, t_2]$, and assume that there are no misreports happening between t_0 and t_2.

In reality, the estimation of the utility loss can be achieved through well-developed regression models based on the available e-marketplace data (e.g., a regression model that associates a seller's revenue with seller reputation based on Taobao transaction data has been specified in [9]). The compensation procedure stop time t_2 can be tried out until the compensated utility is close to the utility loss. Algorithm 1 shows an implementation of the proposed RepRev mechanism. The algorithm is based on the Taobao reputation system and adopted in the experiments .

Procedure: Utility Compensation
Input   : s, seller whose utility being compensated;
               $U(R(t))$, regression model for revenue [9];
               T_s, vector of rating provision time of s;
               P_s, vector of positive ratings of s;
               N_s, vector of negative ratings of s;
               γ_s, the number of misreports;
               t_0, the time the misreports are provided;

1 $\Delta_1 = 0$;
2 **foreach** *transaction time t in T_s and in $[t_0,t_1]$* **do**
3      $R_s(t) = P_s(t) - N_s(t) + 2\gamma_s$;
4      $R'_s(t) = P_s(t) - N_s(t)$;
5      $\Delta_1 = \Delta_1 + U(R_s(t)) - U(R'_s(t))$;
6 $\delta_s = 2\gamma_s$; //according to Eq. (2)
7 $P_s(t_1) = P_s(t_1) + \gamma_s$;
8 $N_s(t_1) = N_s(t_1) - \gamma_s$;
9 $\Delta_2 = 0$;
10 **while** $\Delta_2 < \Delta_1$ **do**
11   **if** *s has a transaction at time t* **then**
12          $R_s(t) = P_s(t) - N_s(t)$;
13          $\Delta_2 = \Delta_2 + U(R_s(t) + \delta_s) - U(R_s(t))$;
14   **if** *s receives a positive rating* **then**
15              $P_s(t) = P_s(t) + 1$;
16   **else**
17              $N_s(t) = N_s(t) + 1$;

Algorithm 1: Utility compensation

In Algorithm 1, we have 7 inputs: 1) the seller s who is waiting to compensate his utility loss; 2) $U(R(t))$, the available regression model for revenue; 3) T_s, the vector storing the rating provision time for the transactions involving s; 4) P_s, the vector storing the total number of positive ratings till a rating provision time specified in T_s; 5) N_s, the vector storing the total number of negative ratings till a rating provision time specified in T_s; 6) γ_s, the number of misreported ratings; 7) t_0, the time when misreported ratings are provided. Lines 2-5 are to estimate the utility loss during the time from t_0 to t_1. More specifically, Lines 3 and 4 are to calculate the seller reputation when there are no misreports, and when there are misreports, respectively. Line 5 is to estimate the utility loss as shown in Figure 1. Line 6 is to calculate the reputation inflation amount. Lines 7 and 8 are to restore seller reputation to the one when there are no misreports. Lines 10-17 are to compensate s's utility loss. More specifically, Line 12 is to inflate s's reputation with δ_s. Line 13 is to calculate the compensated utility during one transaction. Lines 14-17 are to update rating records.

As a summary, RepRev works as follows:

1. After the seller and the buyer agree to revise the seller's reputation, RepRev first restores the seller's reputation to the one when there are no misreports, then it computes the value of δ_s based on the underlying reputation system, and adds it to the seller's restored reputation value.

2. RepRev estimates the utility loss based on the available regression model, and initiates the compensation procedure. During the compensation period, the seller's reputation is inflated by δ_s. The compensation procedure stops at t_2, when the estimated utility loss is fully compensated.

3. After t_2, the seller's reputation is allowed to fluctuate with his behavior again without inflation according to the underlying reputation system.

There are some points worthy noticing. Firstly, we currently focus on a centralized system model which is widely used by existing e-marketplaces. Therefore, the initiation to compensate the seller's utility loss can be controlled by a system arbitrator. Secondly, it is necessary to be careful to initiate the mechanism in order to prevent possible abuse. For example, the seller and the buyer should provide reliable evidence to show that the ratings are really misreported. Thirdly, it is assumed that the number of buyers is much greater than the number of sellers [15]. Thus, the inflation of a seller's reputation increases his revenue, while not significantly affecting other sellers' business.

4. EXPERIMENTAL EVALUATION

In this section, we conduct experiments based on realistic settings to show how the misreported negative ratings impact sellers' utility, and the performance of RepRev in compensating sellers' utility loss.

4.1 Practical Basis of Experimental Settings

Some empirical studies have been conducted to investigate the relationship between sellers' utilities (e.g., price, transaction volume, and revenue) and sellers' reputation [9, 11, 16]. More specifically, a thorough study has been conducted in [9] based on the data collected from Taobao. Regression models have been trained to capture the relationship between sellers' utilities and their reputation. We design our simulations based on the study results from [9].

The reputation system adopted by Taobao is similar to that used by eBay. A seller's reputation is associated with his rating score. A seller may receive a score of 1, 0 and -1 representing a positive, neutral and negative rating for a transaction, respectively. According to the study in [18], neutral ratings almost have the same effect as negative ratings. For simplicity, in our study, we only consider positive

and negative ratings. According to the seller accumulated rating score, there are 21 rating grades as shown in the second and third columns of Table 1[2].

Table 1: Taobao seller rating score, grade and percentage

Index	Score	Grade	Perc.(%)
1	< 4	0	7.35
2	4-10	1 (1 heart)	8.23
3	11-41	2 (2 hearts)	16.18
4	41-90	3 (3 hearts)	11.8
5	91-150	4 (4 hearts)	7.91
6	151-250	5 (5 hearts)	7.84
7	251-500	6 (1 diamond)	11.93
8	501-1,000	7 (2 diamonds)	9.99
9	1,001-2,000	8 (3 diamonds)	7.54
10	2,001-5,000	9 (4 diamonds)	6.31
11	5,001-10,000	10 (5 diamonds)	2.54
12	10,001-20,000	11 (1 blue crown)	1.4
13	20,001-50,000	12 (2 blue crowns)	0.73
14	50,001-100,000	13 (3 blue crowns)	0.16
15	100,001-200,000	14 (4 blue crowns)	0.06
16	200,001-500,000	15 (5 blue crowns)	0.02
17	500,001-1,000,000	16 (1 gold crown)	0
18	1,000,001-2,000,000	17 (2 gold crowns)	0
19	2,000,001-5,000,000	18 (3 gold crowns)	0
20	5,000,001-10,000,000	19 (4 gold crowns)	0
21	> 10,000,000	20 (5 gold crowns)	0

More specifically, the sellers with rating scores lower than 251 are referred to as new sellers, and other sellers are referred to as established sellers. Figure 2 shows an example of the snapshot of a Taobao seller's reputation information[3]. The seller has received 11,904 positive ratings, 1 neutral rating, and 4 negative ratings in the past month. His accumulated rating score is 14,665. His rating grade is 11 (i.e., "1 blue crown"). One more piece of information shown in Figure 2 is the percentage of positive ratings, which is 99.96%.
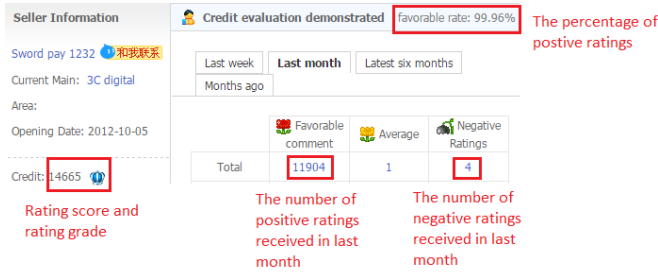

Figure 2: An example of a Taobao seller reputation information

Regression models of the relationship between seller utilities (e.g., revenue, price, and transaction volume) and seller reputation information have been specified in [9] as follows:

$$\ln(U_s(t)) = a_1 G_s(t-1) + a_2 R_s(t-1) + a_3 Q_s(t-1), \quad (7)$$

where $U_s(t)$ is a seller s's utility at time t, $G_s(t-1)$, $R_s(t-1)$, and $Q_s(t-1)$ are s's rating grade, rating score, and posi-

[2]http://service.taobao.com/support/seller/knowledge-847753.htm

[3]http://rate.taobao.com/user-rate-fd16402a5a8f3a087f3440d5cf70080d.htm?spm=2013.1.1000126.3.8nWr4s

tive rating percentage at time $t-1$, respectively. For each aspect of the utility (e.g., revenue, price, transaction volume), different sets of parameter values (i.e., a_1, a_2, and a_3) have been trained. As the rating grade is determined by the rating scores, $U_s(t)$ can be considered as a function of the rating score and the percentage of positive ratings. Furthermore, as the variance of the percentage of positive ratings is small (i.e., mean value of the percentage of positive ratings is 99.467% for the established sellers, and standard deviation is 0.008), we focus on studying the impacts of misreports on the utilities of the sellers with different rating scores.

4.2 Experimental Settings

We simulate 10,000 sellers having transactions following the seller percentage distribution specified in [9], as shown in the fourth column of Table 1. A flowchart of the simulation is illustrated as shown in Figure 3. At the beginning of a simulation, each seller is associated with a randomly generated initial rating score according to his rating grade. For example, there are 7.35% of sellers (i.e, 735 sellers) with rating grade 0. For each of the 735 sellers, an integer value in the range of [0,4) is randomly generated as the seller's initial rating score. For each seller, an initial number of transactions is generated based on the regression model specified in [9] using its initial rating score as the input. The maximum number of transactions among the 10,000 sellers is about 1.544×10^5. Then we simulate that the 10,000 sellers have transactions during 1.544×10^5 time steps. Each seller has an initial probability to have a transaction at a time step, which is simulated by dividing the initial number of transactions for the seller by this maximum transaction number (i.e., 1.544×10^5). Thus, the most popular seller in our simulation has, on average, one transaction per time step. In this way, we can simulate how the sellers' utility is impacted by their rating scores. For example, the higher a seller's rating score is, the more transaction opportunities he will have at any time step, and a higher revenue expected.

At a time step, it is first checked whether a seller will have a transaction according to the probability of the seller having transactions. If the seller is simulated to have a transaction at current time step, it is further checked whether the seller's revenue should be compensated. If his revenue needs to be compensated, the revenue for current transaction is decided by the regression model [9] with the inflated rating score at current time step as input, and the seller's accumulated revenue increases accordingly. Then the seller will receive a binary rating according to the percentage of positive ratings. In our simulations, the compensation procedure is triggered between 1,000 to 2,000 time steps after misreported ratings happen following an independent and identical distribution (i.i.d.).

If the seller's revenue does not need to be compensated, the revenue for current transaction is decided by the regression model [9] with the original rating score at current time step as input, and the seller's accumulated revenue increases accordingly. Then it is checked whether there will be a misreport for current transaction. We simulate that the percentage of misreported ratings is varied from 1‰ to 10‰ with 1‰ increment to simulate the occasional happening of the misreported ratings. The occurrence of misreported ratings are uniformly distributed over the time steps which are not in the compensation period in each simulation. If current transaction is simulated to be misreported,
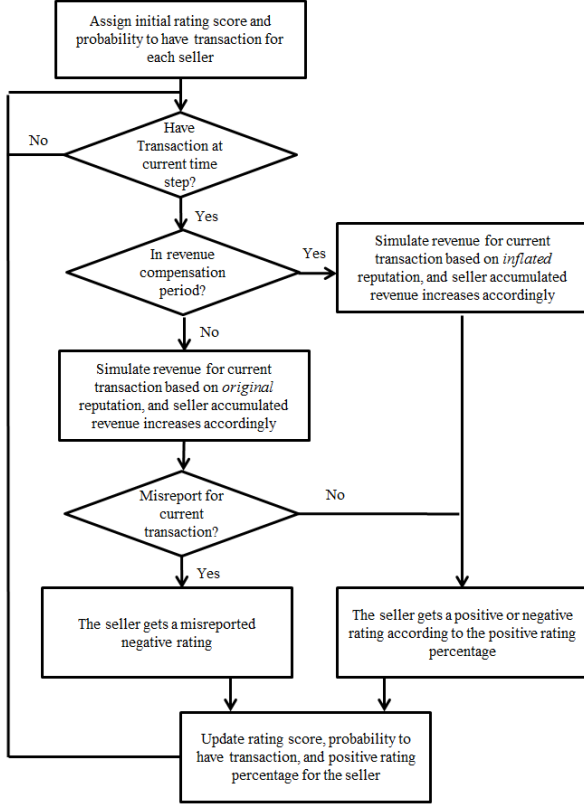
Figure 3: The average percentage of revenue loss due to misreports for sellers with different rating grades

the seller will receive a negative rating if he actually deserves a positive rating. Otherwise, the seller will receive a binary rating which is decided by the the percentage of positive ratings.

Finally, the seller's rating score, probability to have transactions, and percentage of positive ratings are updated. The simulation will end when the maximum time step (i.e., 1.544×10^5) is reached. We conduct each simulation 100 times to improve the statistical accuracy.

4.3 Impacts of Misreported Negative Ratings

We first study the effects of the misreported negative ratings on the accumulated revenue of the sellers. Due to the rare opportunities of the new sellers to have transactions (i.e., only 1 or 2 transactions are simulated in the 1.544×10^5 time steps), we focus on studying the impacts of misreported ratings on the established sellers.

We specifically study the negative effects of misreported ratings on one of the most serious concerns of the sellers – the accumulated revenue. The accumulated revenue $A_s(T)$ for a seller s until time step T is expressed as:

$$A_s(T) = \sum_{t=1}^{T} H_t U_s(t), \qquad (8)$$

where $H_t = 1$ if s has a transaction at time step t, and 0 otherwise. $U_s(t)$ is the revenue for one transaction at time step t.

Figure 4(a) shows the accumulated revenue at the end of

the simulation for the sellers whose initial rating scores are in the range of 251 to 2,000. This group of sellers constitutes 72.4% of the established sellers. From Figure 4(a), it can be seen that there is a trend that the accumulated revenue increases with a seller's rating score.

More specifically, we study the impacts of the misreports on revenue through the *revenue difference ratio* metric (η_s) for each seller s, which is calculated as follows:

$$\eta_s = \frac{A_s(T) - \hat{A}_s(T)}{A_s(T)}, \qquad (9)$$

where $\hat{A}_s(T)$ and $A_s(T)$ are the actual revenue received by a seller s with and without misreports until time T, respectively. A positive η_s value indicates that a seller s's revenue drops with misreported ratings, whereas a negative η_s value means s's revenue increases with misreported ratings. Under this metric, Figure 4(b) shows the percentage of the established sellers whose revenues have been negatively affected by misreports under various misreport percentage settings (i.e., the sellers with $\eta_s > 0$). It can be observed that over half of the established sellers are affected, and as the misreport percentage increases, the percentage of affected sellers also increases, reaching over 80% eventually.

Figure 4(c) shows the average percentage of revenue loss due to misreports. The result presents a similar trend to the percentage of sellers affected by misreports as shown in Figure 4(b). With the misreport percentage increasing, the revenue loss also increases.

Figure 4(d) shows the average η_s values for the established sellers with different rating grades (from 6 to 15). It can be seen that the sellers with lower rating grades (i.e., those with fewer positive ratings either due to fewer number of past transactions or poorer performance in the past) tend to be more significantly affected by misreports than those with higher rating grades. As many sellers with low rating grades are those who have recently started their e-commerce business and in the process of building up their reputation, it is very important for this group of sellers to be protected from the negative effects of misreports so as to enable the e-marketplace as a whole to grow healthily.

4.4 Effectiveness of the RepRev Mechanism

In this part, we present the results of the proposed RepRev mechanism in compensating sellers' utility loss, where the compensation procedure follows Algorithm 1. To study the effectiveness of the proposed RepRev mechanism, we adopt the metric of *revenue difference ratio after compensation* (ξ_s) for each seller s, which is expressed as follows:

$$\xi_s = \frac{A_s(T) - \tilde{A}_s(T)}{A_s(T)}, \qquad (10)$$

where $\tilde{A}_s(T)$ is s's accumulated revenue until T with compensation provided by the proposed RepRev mechanism. The closer the ξ_s value is to 0, the more accurate the provided compensation is.

Figure 4(e) shows the percentage of the sellers negatively affected by misreports after being compensated by RepRev (i.e., the sellers with $\xi_s > 0$). It can be seen that the percentage of the sellers negatively affected by misreports has become stable and does not increase with misreport percentage. The average revenue loss for all the established sellers due to misreports after being mitigated by RepRev has also

Figure 4: The impacts of misreports and the effectiveness of RepRev: (a) Seller revenue changes with rating score; (b) The percentage of sellers affected by misreports; (c) The average percentage of revenue loss due to misreports; (d) The average percentage of revenue loss due to misreports for sellers with different rating grades; (e) The percentage of sellers affected by misreports after compensation; (f) The average percentage of revenue loss due to misreports after compensation; (g) The average percentage of revenue difference before and after compensation when there are low-level misreports; (h) The average percentage of revenue difference before and after compensation when there are there are medium-level misreports; (i) The average percentage of revenue difference before and after compensation when there are high-level misreports.

stabilized to around 1% of the total accumulated revenue as shown Figure 4(f). Compared to Figure 4(c), RepRev has reduced the negative effects of misreports on sellers' revenue by over 50% on average.

Figure 4(g), 4(h) and 4(i) take a more fine grained view on the effectiveness of RepRev for sellers with different rating grades under different levels of misreports. Under *low* level of misreports (with misreport percentage less than 4‰), the effectiveness of RepRev for the established sellers is not significant as shown in Figure 4(g). This is because the negative effects of misreports under this setting is already very low.

Under *medium* level of misreports (with misreport percentage between 4‰ and 6‰), the benefit of RepRev for sellers with rating grades from 6 to 9 has become significant as shown in Figure 4(h). Under *high* level of misreports (with misreport percentage between 6‰ and 10‰), the benefit of RepRev for sellers with rating grades from 6 to 13 has become significant as shown in Figure 4(i).

In summary, the experimental results suggest that RepRev can effectively compensate the sellers' utility loss due to misreports, especially for those within low to medium rating grades (i.e, grades 6-13), who constitute 99.8% of the es-

tablished sellers, and are also the ones that most need their reputation restored and utility recovered.

5. CONCLUSIONS AND FUTURE WORK

In this paper, we have proposed a reputation revision mechanism to compensate the sellers' utility loss when they suffer from misreports which are caused by some unintentional factors such as miscoummunications. The proposed mechanism serves as a more principled "undo" function compared to the existing approaches employed by existing e-marketplaces (e.g., eBay and Taobao) for buyers to ratify mistakes they have made when judging a seller. The proposed mechanism temporarily inflates the reputation of the misjudged seller for a period of time to regain expected utility loss based on the empirical evidence that utility increases with reputation. Simulations based on realistic settings demonstrate the adverse effects of negative misreported ratings on the revenue of the sellers, and the effectiveness of the proposed mechanism in mitigating these impacts, especially for sellers with low and medium rating grades.

In the future, we will investigate more complex scenarios of compensating utility loss due to misreports (e.g., more misreports occurring during the period of compensation, and taking changes in a seller's quality of service into account when compensating their reputation) so as to improve the practical applicability of the proposed mechanism.

6. ACKNOWLEDGMENTS

7. REFERENCES

[1] http://bbs.taobao.com/catalog/thread/154521-259744203.htm.

[2] http://community.ebay.com/t5/archive-feedback/recall-a-feedback/qaq-p/5865187.

[3] http://pages.ebay.com.sg/help/feedback/revision-request.html.

[4] http://wenwen.sogou.com/z/q156781467.htm.

[5] http://zhidao.baidu.com/question/241259999.html.

[6] www.ebay.com.

[7] www.taobao.com.

[8] C. Dellarocas. Immunizing online reputation reporting systems against unfair ratings and discriminatory behavior. In *Proceedings of the 2nd ACM Conference on Electronic Commerce*, pages 150–157, 2000.

[9] Y. Fan, J. Ju, and M. Xiao. Reputation premium and reputation management: Evidence form the largest e-commerce platform in china. *Working paper*, 2014.

[10] H. Fang, J. Zhang, and N. M. Thalmann. Subjectivity grouping: Learning from users' rating behavior. In *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems*, pages 1241–1248, 2014.

[11] G. Jolivet, B. Jullien, and F. Postel-Vinay. Reputation and prices on the e-market: Evidence from a major french platform. *Working Paper*, 2013.

[12] A. Jøsang and R. Ismail. The beta reputation system. In *Proceedings of the 15th Bled Electronic Commerce Conference*, 2002.

[13] S. Liu, A. C. Kot, C. Miao, and Y.-L. Theng. A dempster-shafer theory based witness trustworthiness model to cope with unfair ratings in e-marketplace. In *Proceedings of the 14th Annual International Conference on Electronic Commerce*, pages 99–106, 2012.

[14] S. Liu, H. Yu, C. Miao, A. C. Kot, and Y.-L. Theng. A fuzzy logic based reputation model against unfair ratings. In *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems*, pages 821–828, 2013.

[15] P. J. McNulty. A note on the history of perfect competition. *Journal of Political Economy*, 75(4):395–399, 1967.

[16] Y. Qiang, M. Xu, M. Kianga, W. Wu, and F. Sun. In-depth analysis of the seller reputation and price premium relationship: A comparison between ebay us and taobao china. *Journal of Electronic Commerce Research*, 14(1):1–10, 2013.

[17] K. Regan, P. Poupart, and R. Cohen. Bayesian reputation modeling in e-marketplaces sensitive to subjectivity, deception and change. In *Proceedings of the 21st National Conference on Artificial Intelligence*, pages 206–212, 2006.

[18] P. Resnick and R. Zeckhauser. Trust among strangers in internet transactions: Empirical analysis of ebay's reputation system. *The Economics of the Internet and E-Commerce*, 11(2):127–157, 2002.

[19] W. T. L. Teacy, N. R. Jennings, N. R. Rogers, and M. Luck. A hierarchical bayesian trust model based on reputation and group behaviour. In *the 6th European Workshop on Multi-Agent Systems*, pages 206–212, 2008.

[20] W. T. L. Teacy, J. Patel, N. R. Jennings, and M. Luck. TRAVOS: trust and reputation in the context of inaccurate information soucres. *Autonomous Agents and Multi-Agent Systems*, 12(2):183–198, 2006.

[21] J. Weng, C. Miao, and A. Goh. An entropy-based approach to protecting rating systems from unfair testimonies. *Transactions on Information and Systems*, E89-D(9):2502–2511, September 2006.

[22] A. Whitby, A. Jøsang, and J. Indulska. Filtering out unfair ratings in bayesian reputation systems. *ICFAIN Journal of Management Research*, 4(2):48–64, 2005.

[23] H. Yu, Z. Shen, and B. An. An adaptive witness selection method for reputation-based trust models. In *Proceedings of the 15th International Conference on Principles and Practice of Multi-Agent Systems (PRIMA)*, pages 184–198, 2012.

[24] H. Yu, Z. Shen, C. Leung, C. Miao, and V. R. Lesser. A survey of multi-agent trust management systems. *IEEE Access*, 1(1):35–50, 2013.

[25] H. Yu, Z. Shen, C. Miao, C. Leung, and D. Niyato. A survey of trust and reputation management systems in wireless communications. *Proceedings of the IEEE*, 98:1755–1772, 2010.

[26] J. Zhang and R. Cohen. Evaluating the trustworthiness of advice about selling agents in e-marketplaces: A personalized approach. *Electronic Commerce Research and Applications*, 7(3):330–340, 2008.