Evaluation and Design of Online Cooperative Feedback Mechanisms for Reputation Management

Ming Fan, Yong Tan, and Andrew B. Whinston

Abstract—This research evaluates and analyzes existing mechanisms of online reputation systems based on cooperative feedbacks of past transaction information. We find existing popular feedback systems do not provide sustained incentives for sellers to behave honestly over time. We propose a new design of reputation system based on exponential smoothing. This mechanism is shown to be more robust compared to the existing systems. We relax the assumption of a static product value and have implemented a two-level exponential smoothing policy to process and aggregate reputation information. The simulation results show that the policy can serve as a sustained incentive mechanism for the seller.

Index Terms—Electronic commerce, online trust, reputation systems, exponential smoothing.

1 Introduction

With the enormous growth of the Internet and electronic commerce, online trust has become an increasingly important issue. Although the Internet offers tremendous new opportunities in business and other areas, there are great uncertainties and risks in the online world. Over the Internet, agents[1] often have to deal with other agents who live far away and can only be identified by an e-mail address. This requires a great amount of trust given that it may not be possible to track down or even identify the other party. However, trust is difficult to build among strangers over the network because agents lack past history information about each other. Given these factors, the temptation to defect could outweigh the incentive to cooperate. Indeed, cheating and fraud are real over the Internet. It was reported that fraud on the Internet rose sharply in 2002, with online auction fraud accounting about half of all the complaints [21].

The issues of trust and trust creation have been an active research area in different disciplines including economics [11], [18], sociology [5], business information systems [2], [3], [17], and multiagent systems [4], [19], [23]. From an economics perspective, agents will not cooperate when the interaction lasts for only one period, as indicated in the results of the well-known prisoner's dilemma game. However, cooperation can be sustained when the game is played repeatedly, even with a finite horizon [14]. In a multiagent system, agents are assumed to be self-interested and each agent's behavior is determined by its calculated long-term payoff. This is consistent with the game theoretical approach employed in economics literature. Trust among agents primarily involves a calculative process as agents figure out the costs of cheating and the rewards of staying honest. To the extent that the benefits of cheating do not exceed the costs, the agent would infer that it is not in his best interest to cheat and, therefore, the agent can be trusted.

There has been an exciting development in online practices to promote trust, among which reputation systems have been catching the most attention. For example, eBay, the largest person-to-person online auction site, has deployed a reputation system. Although eBay only offers limited insurance and buyers and sellers face significant risks, it seems that most of the buyers and sellers trust each other in completing the promised transactions. eBay attributed this to its reputation system [17]. At eBay, after a transaction is complete, the buyer has the opportunity to rate the seller and leave comments. A numerical rating is associated with the comment, indicating whether the comment is positive (+1), negative (-1), or neutral (0). To calculate the aggregate reputation score for a seller, eBay aggregates the reputation scores over time. The reputation system serves as both a source of information for new buyers and potential means of sanctions for sellers who committed bad behavior. Reputation systems also exist in other online businesses. For example, Amazon.com and Yahoo! implement their own reputation systems for their online auctions. Bizrate.com rates online retailers by asking consumers to complete a survey after each purchase. Epinions.com reviews products and services ranging from computers to restaurants. Reputations can also be applied in software agent systems [4] and for locating Web services.

1. We use agent in the paper to refer to a decision-making unit. It could be an individual or even an automated software application with decision rules.

- *M. Fan and Y. Tan are with the University of Washington Business School, PO Box 353200, Seattle, WA 98195-3968. E-mail: {mfan, ytan}@u.washington.edu.*
- *A.B. Whinston is with Center for Research in Electronic Commerce, University of Texas at Austin, Austin, TX 78712. E-mail: abw@uts.cc.utexas.edu.*

Manuscript received 30 June 2003; revised 5 Feb. 2004; accepted 23 June 2004; published online 17 Dec. 2004.
For information on obtaining reprints of this article, please send e-mail to: tkde@computer.org, and reference IEEECS Log Number TKDE-0106-0603.

Essentially, a reputation system is a cooperated knowledge sharing mechanism among market participants. In typical models of repeated games, a small number of agents interact with each other repetitively. Therefore, each interaction experience is perfectly observed and remembered by the players. However, this is not the case in the real world. First, there are a large number of agents. In addition, although transactions are repetitive, an agent will likely interact with a different set of agents in the next period. Thus, we cannot expect an agent who was cheated earlier to be able to punish the cheater directly. Cooperative knowledge sharing of past transaction information among market participants is necessary in order to enforce community norms and punish fraud. A dishonest agent's behavior needs to be known to the community and the agent has to be sanctioned by other agents in the community.

According to eBay, only 0.01 percent of all completed auctions are fraudulent [9]. But, it is widely believed the real number is higher since the reported number is drawn from people who go through eBay's fraud-insurance claim process. Many people for some reason did not complete the process. It is estimated that, in the computer and consumer electronics auctions, theft represents as much as 5 percent of all auctions [9]. It appears that the effectiveness of online reputation systems is not all that clear in both theory and practice. Prior empirical studies have shown that the eBay reputation system has a small but statistically significant positive effect on selling price [2], [12], [15], [16]. However, there are many questions with the eBay-type reputation mechanism. In [22], empirical results show that eBay's system does not provide sustained incentives for reputable sellers to provide honest quality over time. There are also possibilities that a seller could change his identity and use pseudonyms [10]. In [7], new mechanisms have been proposed to reduce the negative effects of both unfairly high and low ratings.

The success of online reputation systems depends on many factors. First, it relies on the active participation of the agents in the community [8] and unbiased feedbacks from those participants. Second, the shared information has to be processed and presented in the most meaningful format. At present, eBay aggregates the reputation scores while Amazon.com calculates the average rating score as indicators of the trustworthiness of the sellers. Different reputation mechanisms have strong implications on the trust behavior of the buyers and sellers. In this research, we aim to design an alternative reputation mechanism that can provide incentives for sellers to stay honest over time. One big challenge in designing such a system is that agents usually change auction items over time. Potentially, an agent can develop a high reputation by selling small-value items and may decide to cheat in selling an expensive product later on. Therefore, the design of the mechanism has to be robust enough to adjust the seller agent's rating accordingly in order to deter cheating.

This research is broadly related to mechanism design literature in economics, which concerns the development of institutional rules that can be created and adopted by market makers to govern market transactions [6]. Following the design rules, the dominant strategy for an agent is to take actions that implement the desired social choice. In [11], the institutional structure of the medieval merchant guilds was studied from a game theoretical approach. This research is also related to credit and risk rating in financial markets [13].

The rest of the paper is organized as follows: In the next section, we analyze and evaluate accumulative and average reputation mechanisms. We propose a reputation mechanism design based on exponential smoothing in Section 3. We relax the assumption of a constant auction item value and present numerical and simulation results in Section 4. Finally, we conclude the paper and discuss limitations and future research issues in Section 5.

2 Evaluating Accumulative and Average Mechanisms

In this section, we develop a model to evaluate two reputation mechanisms that are common in the online environment: accumulative score mechanism and average score mechanism. The purpose is to examine whether those mechanisms provide sellers incentives to consistently provide quality products as advertised over time. We analyze a repeated auction of an identical good from a single seller, who maximizes his profit and will provide an inferior product given opportunities. In Section 4, the identical good assumption will be relaxed. We also assume that buyers cannot verify the product quality ex ante. They can only examine the quality after they receive the delivery. Different from standard settings in repeated games, we examine a repeated game with the same seller but a different set of bidders over time. However, we assume the number of bidders is fixed over time and bidder i's valuation of the item v_i follows a distribution F. Since the bidders change over time, they are unable to observe the seller's past behavior directly. They can only observe the feedback left by previous buyers. The feedback score s_t that the seller gets is based on the buyer's feedback for the transaction that occurred at time t.

The following is the sequence of our analytical model: 1) The market maker implements a reputation mechanism. 2) Bidders make bids on the auction according to the advertised product specification and the seller's reputation information. 3) The seller decides the quality of the product to ship to the winning bidder and receives a feedback rating from the buyer.

2.1 Accumulative Score Mechanism

We denote R_t^a as the accumulated reputation score at t. R_t^a is a simple summation of the seller's feedback scores s_i in all previous periods, i.e., $R_t^a = \sum_{i=1}^{t-1} s_i$. This mechanism is currently used by eBay. Here, we assume that buyers provide unbiased feedbacks. As mentioned earlier, the value of s_t could be 1 (positive feedback), -1 (negative feedback), or 0 (neutral feedback). We let N_t denote the total number of negative ratings at time t. Since the number of negative ratings is also provided by eBay, we assume bidders will include that factor in their bidding function as well.

Most of the online auctions allow proxy bidding, which is equivalent to the second-price auction. Thus, a bidder has the incentive to bid his truthful valuations, denoted as v_i [12], [20]. In an environment with uncertain seller quality, a bidder will likely bid a risk-adjusted valuation based on the seller's reputation and his true valuation. Here, we use the following bid function to represent the seller's belief of the bidders' average bidding behavior:

$$b_i = B(R_t^a, N_t)v_i,$$

where $B(R_t^a, N_t)$ represents the discount of transaction risk based on the seller's reputation. This approach was proposed and supported empirically in [12]. In addition, we note that the bid function $B(R_t^a, N_t) \leq 1$ since bidders will not bid more than their valuations.

In equilibrium, the transaction price for the auction is $B(R_t^a, N_t)v_{(2)}$, where $v_{(2)}$ is the expected second highest valuation among all the bidders. To simplify the notation, we use v to represent $v_{(2)}$ from now on. Thus, the equilibrium transaction price becomes $B(R_t^a, N_t)v$.

With an accumulative reputation system, the bidding function is increasing and concave on the reputation score, i.e., $\partial B/\partial R_t^a > 0$ and $\partial^2 B/(\partial R_t^a)^2 < 0$ [22], and decreasing and concave on the number of negative ratings, i.e., $\partial B/\partial N_t < 0$ and $\partial^2 B/(\partial N_t)^2 < 0$. Empirical studies have shown that negative comments negatively impact the bidding value [2]. When negative comments are few, bidders might have a certain tolerance toward the seller given the possibility of rating errors and disputes. However, when negative comments increase, we expect the impacts of those ratings will have an accelerating effect. In addition, we assume the effect of negative ratings does not increase with reputation scores, i.e. $\partial^2 B/(\partial R_t^a \partial N_t) \geq 0$. It implies that the bidders' tolerance toward negative ratings increases as the accumulative reputation score goes up. This is easy to understand because the ratio of negative comments to the total number of transactions is lower with a higher reputation score.

Let L_t be the total number of transactions, then we have $L_t = R_t^a + 2N_t$. Therefore, with the introduction of negative ratings, we can see that, given the same reputation score R_t^a, the more transactions the seller has in the past, the more severely the bidder will discount that reputation score since $\partial B/\partial L_t = \partial B/\partial N_t \cdot \partial N_t/\partial L_t < 0$.

Now, we examine the seller behavior under the accumulative mechanism. The seller's profit function at period t is

$$\pi_t = B(R_t^a, N_t)v - \frac{1+s_t}{2}c_1 - \frac{1-s_t}{2}c_0,$$

where c_1 is the cost if the seller behaves honestly and provides the quality as advertised. c_0 is the cost if the seller provides inferior quality. We have $v > c_1 > c_0 > 0$. Here, we use the feedback score from the buyer, $s_t (s_t = 1 \text{ or } -1)$, to represent the decision for the seller. Under the assumption that buyers provide unbiased feedback, $s_t = 1$ indicates that the seller provided honest quality at t and $s_t = -1$ indicates otherwise. The total profits for the seller from the current period t are

$$\Pi_t = \sum_{k=t}^{\infty} \delta^{k-t}\left(B(R_k^a, N_k)v - \frac{1+s_k}{2}c_1 - \frac{1-s_k}{2}c_0\right),$$

where δ is the discount factor.

We adopt the dynamic programming approach to examine the seller's profit. The above problem can be rewritten iteratively as

$$\Pi_t(R_t^a, N_t, s_t) = B(R_t^a, N_t)v - \frac{1+s_t}{2}c_1 - \frac{1-s_t}{2}c_0 + \delta \cdot \hat{\Pi}_{t+1}(R_{t+1}^a, N_{t+1}), \quad (1)$$

where

$$\hat{\Pi}_t(R_t^a, N_t) = \max_{s_t} \Pi_t(R_t^a, N_t, s_t) \quad (2)$$

is the optimal discounted profit starting from the period t with a reputation score of R_t^a and number of negative ratings of N_t. Entering period t with a reputation score R_t^a, the seller chooses s_t to not only maximize the profit in the current period, but also take into account the impact of s_t on future streams of profits through the updated reputation score R_{t+1}^a and the number of negative ratings of N_{t+1}. Although s_t is discrete, we use a continuous function to approximate it since the function is well behaved. Taking a derivative of (1) with respect to s_t, we get

$$\frac{\partial \Pi_t}{\partial s_t} = -\frac{c_1 - c_0}{2} + \delta\left(\frac{\partial \hat{\Pi}_{t+1}}{\partial R_{t+1}^a} - \frac{1}{2}\frac{\partial \hat{\Pi}_{t+1}}{\partial N_{t+1}}\right). \quad (3)$$

Here, $\partial R_{t+1}^a/\partial s_t = 1$ and $\partial N_{t+1}/\partial s_t = -1/2$ since we can express $R_{t+1}^a = R_t^a + s_t$ and $N_{t+1} = N_t + (1-s_t)/2$ for the accumulative reputation mechanism. In our setting, s_t can only have two values, therefore, if $\partial \Pi_t/\partial s_t > 0$, $s_t = 1$, while, if $\partial \Pi_t/\partial s_t < 0$, $s_t = -1$.

To further simplify $\partial \Pi_t/\partial s_t$, we derive $\partial \hat{\Pi}_{t+1}/\partial R_{t+1}^a$ and $\partial \hat{\Pi}_{t+1}/\partial N_{t+1}$ using (1) and (2):

$$\frac{\partial \hat{\Pi}_{t+1}}{\partial R_{t+1}^a} = v\frac{\partial B}{\partial R_{t+1}^a} + \delta \frac{\partial \hat{\Pi}_{t+2}}{\partial R_{t+2}^a}\frac{\partial R_{t+2}^a}{\partial R_{t+1}^a} = v\sum_{k=t}^{\infty}\delta^{k-t}\frac{\partial B(R_{k+1}^a, N_{k+1})}{\partial R_{k+1}^a},$$

$$\frac{\partial \hat{\Pi}_{t+1}}{\partial N_{t+1}} = v\frac{\partial B}{\partial N_{t+1}} + \delta \frac{\partial \hat{\Pi}_{t+2}}{\partial N_{t+2}}\frac{\partial N_{t+2}}{\partial N_{t+1}} = v\sum_{k=t}^{\infty}\delta^{k-t}\frac{\partial B(R_{k+1}^a, N_{k+1})}{\partial N_{k+1}}.$$

Substituting the above two equations into (3), we have,

$$\frac{\partial \Pi_t}{\partial s_t} = -\frac{c_1 - c_0}{2} + \delta v \sum_{k=t}^{\infty}\delta^{k-t}\left(\frac{\partial B(R_{k+1}^a, N_{k+1})}{\partial R_{k+1}^a} - \frac{1}{2}\frac{\partial B(R_{k+1}^a, N_{k+1})}{\partial N_{k+1}}\right). \quad (4)$$

We explore the possibility that, when the reputation score is large enough, the dishonest seller just cheats once and then remains honest afterward. Since $\partial B(R_{t+1}^a, N_{t+1})/\partial R_{t+1}^a$ decreases in R_{t+1}^a and $-\partial B(R_{t+1}^a, N_{t+1})/\partial N_{t+1}$ does not increase in R_{t+1}^a, we have

$$\frac{\partial \Pi_t}{\partial s_t} < -\frac{c_1 - c_0}{2} + \frac{\delta v}{1-\delta}\left(\frac{\partial B(R^a_{t+1}, N_{t+1})}{\partial R^a_{t+1}} - \frac{1}{2}\frac{\partial B(R^a_{t+1}, N_{t+1})}{\partial N_{t+1}}\right). \tag{5}$$

The right-hand side of (5) is a decreasing function of R^a_{t+1} and $\partial B/\partial R^a$ will approach zero with a high enough R^a because the bidding function is concave and bounded. Therefore, the right-hand side will become negative as R^a_{t+1} is large enough.

Similarly, we can find an upper bound N^* for the number of negative ratings for a given reputation score. From (4), we can derive

$$\frac{\partial \Pi_t}{\partial s_t} > -\frac{c_1 - c_0}{2} - \frac{\delta v}{2}\frac{\partial B(R^a_{t+1}, N_{t+1})}{\partial N_{t+1}}.$$

Setting the right-hand side of the above inequality to zero, we can find the $N^*(R^a)$ beyond which it is not in the seller's interest to behave dishonestly. It can be shown, using the properties of the bid function, that the solution $N^*(R^a)$ exists and it increases with the reputation score.

Proposition 1. *Under the accumulative score mechanism and an infinite horizon, always being honest is not the best strategy for the seller, i.e., $\{s^*_t = 1, \forall t\}$ is not possible. There also exists an upper bound on the number of cheatings a seller can conduct for a given reputation score.*

Equation (5) sets an upper bound for $\partial \Pi_t / \partial s_t$. The $(c_1 - c_0)/2$ represents part of the gains of cheating for the seller, while the second term of the right-hand side of (5) is the upper bound of future gains from current honest behavior. When R^a_t increases and becomes very large, $\partial B(R^a_{t+1})/\partial R^a_{t+1}$, the gain from a higher reputation, will diminish. Therefore, it is in the seller's interest to cheat at least once. Meanwhile, the disclosure of the total number of negative scores forbids the seller from getting away if he cheats repeatedly. Overall, the accumulative mechanism is less perfect since the accumulative reputation score provides a weak incentive for highly reputable sellers.

2.2 Average Score Mechanism

Another popular reputation mechanism computes average reputation scores based on feedbacks. Using the average score mechanism, the seller's reputation score at t is the average of the feedback scores in the previous $t-1$ periods, explicitly,

$$R^m_t = \frac{1}{t-1}\sum_{i=1}^{t-1} s_{t-i}.$$

After the rating for period t is posted, the seller's reputation average at $t+1$ will be updated according to the following iterative relation,

$$R^m_{t+1} = \frac{1}{t}s_t + \frac{t-1}{t}R^m_t.$$

Technically, $R^m_t \in [-1, 1]$. However, since a seller can rarely stay in business when his reputation becomes negative, we can safely assume $R^m_t \in [0, 1]$. Since R^m_t is bounded, we assume a simple linear bidding function and let $b_i = R^m_t v_i$, instead of using a concave bidding function employed in accumulative score mechanism. In equilibrium, the transaction price for the auction is $R^m_t v$. Therefore, at period t, the seller's profit is

$$\pi_t(R^m_t, s_t) = R^m_t v - \frac{1+s_t}{2}c_1 - \frac{1-s_t}{2}c_0. \tag{6}$$

Similarly, the seller's dynamic programming formulation is

$$\hat{\Pi}_t(R^m_t) \equiv \max_{s_t} \Pi_t(R^m_t, s_t) = \max_{s_t}\left(R^m_t v - \frac{1+s_t}{2}c_1 - \frac{1-s_t}{2}c_0 + \delta \cdot \hat{\Pi}_{t+1}(R^m_{t+1})\right). \tag{7}$$

The derivative of (7) with respect to s_t is

$$\frac{\partial \Pi_t}{\partial s_t} = -\frac{c_1 - c_0}{2} + \frac{\delta}{t}\frac{\partial \hat{\Pi}_{t+1}}{\partial R^m_{t+1}}. \tag{8}$$

Evaluating (7) at $t+1$ and taking a derivative of $\hat{\Pi}_{t+1}(R^m_{t+1})$ with respect to R^m_{t+1}, we get,

$$\frac{\partial \hat{\Pi}_{t+1}}{\partial R^m_{t+1}} = v + \delta \frac{\partial \hat{\Pi}_{t+2}}{\partial R^m_{t+2}}\frac{\partial R^m_{t+2}}{\partial R^m_{t+1}} = v + vt\sum_{n=1}^{\infty}\frac{\delta^n}{t+n}. \tag{9}$$

Substituting (9) into (8), we have

$$\frac{\partial \Pi_t}{\partial s_t} = -\frac{c_1 - c_0}{2} + \delta v \sum_{n=0}^{\infty}\frac{\delta^n}{t+n}. \tag{10}$$

In (10), $(c_1 - c_0)/2$ represents the gains of cheating for the seller, while the second term $\delta v \sum_{n=0}^{\infty} \delta^n/(t+n)$ is the accumulative future benefits of the current honest behavior. Similar to the accumulative score mechanism, when the long-term benefit of a good reputation outweighs the initial opportunistic behavior, the seller will choose to provide honest quality. We can also easily observe that $\partial_\delta(\partial \Pi_t/\partial s_t) > 0$. This indicates that the future benefits of a good reputation increase with δ, the discount factor.

Here, we consider the case that the seller at least provides honest quality service at the beginning. This requires (10) to be positive for small values of t. If we set $t = 1$, (10) becomes,

$$\left.\frac{\partial \Pi_t}{\partial s_t}\right|_{t=1} = -\frac{c_1 - c_0}{2} - v\ln(1-\delta).$$

This defines a lower bound for the discount rate, $\delta > 1 - e^{-(c_1-c_0)/2v}$, which forces the seller to behave honestly, at least in the first period. Since $v > c_1 > c_0 > 0$, we find that $(c_1 - c_0)/2v < 1/2$. Thus, as long as $\delta > 0.4$, the condition for the seller to provide honest quality in the first period is satisfied.

It is straightforward to verify that the right-hand side of (10) is a strictly decreasing function of t. We define t^* as the threshold value of t beyond which the right-hand side of (10) becomes negative. That is, for $t > t^*$, $\partial \Pi_t / \partial s_t < 0$ and the seller will choose to provide inferior quality.

Proposition 2. *Under the average score mechanism and an infinite horizon, when $t < t^*$, the seller's dominant strategy is* *to provide honest quality; when $t > t^*$, the seller's dominant strategy is to provide inferior quality, where t^* satisfies*

$$-\frac{c_1 - c_0}{2} + \delta v \sum_{n=0}^{\infty} \frac{\delta^n}{t^* + n} = 0.$$

The fact that (10) is a decreasing function of t indicates the marginal benefit of staying honest decreases over time. The explanation for that is, when the history is long enough, the current behavior has almost no effect on the future reputation scores under the average score mechanism. Therefore, over time, the incentive to stay honest weakens for the seller.

2.3 Seller Identity Change

The analysis, thus far, assumes that the seller has a perception of an infinite horizon. We now examine the possibility that the seller has an option to change his identity. Over the Internet, people can stop using the current identity and create a new account with virtually no cost [10]. Here, we examine a scenario that a seller decides whether to stay honest or not at time t. Γ is the future date the seller considers a change of identity. Although the timing of Γ itself is an interesting problem, it is beyond the scope of this study. We examine the seller's behavior when Γ is given.

First, we examine the game right before the seller plans to change its identity. In that last stage, we have the seller's payoff as

$$\Pi_\Gamma = Bv - \frac{1 + s_\Gamma}{2} c_1 - \frac{1 - s_\Gamma}{2} c_0.$$

Taking the first derivative, we have

$$\frac{\partial \Pi_\Gamma}{\partial s_\Gamma} = -\frac{c_1 - c_0}{2} < 0,$$

for both accumulative and average mechanisms. Therefore, the dominant strategy is for the seller to cheat in the final stage, regardless of the different reputation mechanisms.

Since only the seller knows the timeline of the game, but not the potential bidders, bidders will continue to bid according to the bid function. Using the accumulative score mechanism, the seller's dynamic programming problem is similar to (1). The only difference is that, at time t, there are $\Gamma - t$ stages toward the end. The derivative of the objective function with respect to s_t becomes

$$\frac{\partial \Pi_t}{\partial s_t} = -\frac{c_1 - c_0}{2} + \delta v \sum_{k=t}^{\Gamma-1} \delta^{k-t} \left(\frac{\partial B(R^a_{k+1}, N_{k+1})}{\partial R^a_{k+1}} - \frac{1}{2} \frac{\partial B(R^a_{k+1}, N_{k+1})}{\partial N_{k+1}} \right).$$

Similarly, we derive the upper bound for $\partial \Pi_t / \partial s_t$,

$$\frac{\partial \Pi_t}{\partial s_t} < -\frac{c_1 - c_0}{2} + \delta v \frac{1 - \delta^{\Gamma - t}}{1 - \delta} \left(\frac{\partial B(R^a_{t+1}, N_{t+1})}{\partial R^a_{t+1}} - \frac{1}{2} \frac{\partial B(R^a_{t+1}, N_{t+1})}{\partial N_{t+1}} \right). \quad (11)$$

The upper bound for $\partial \Pi_t / \partial s_t$ in a finite horizon specified in (11) is lower than the upper bound of $\partial \Pi_t / \partial s_t$ given by (5) for the infinite horizon case.

Proposition 3. *Under the accumulative mechanism, it is more likely for the seller to behave dishonestly when he can change his identity at Γ, compared to the infinite horizon case.*

When either the reputation score is high enough or it is getting close to Γ, the seller has higher incentive to cheat. Similarly, in evaluating the average score mechanism when the seller changes the identity at Γ, we have

$$\frac{\partial \Pi_t}{\partial s_t} = -\frac{c_1 - c_0}{2} + \delta v \sum_{k=t}^{\Gamma-1} \frac{\delta^{k-t}}{k}. \quad (12)$$

Obviously, the seller's decision depends both on both t and $\Gamma - t$. Either a high t, which implies a long history, or that t is getting closer to Γ can lead to $\partial \Pi_t / \partial s_t < 0$.

Proposition 4. *Under the average score mechanism, the seller is more likely to behave dishonestly when he can change his identity at Γ, compared to the infinite horizon case.*

Since $\partial \Pi_t / \partial s_t$ as specified in (12) is always smaller than its corresponding term in (10), $\partial \Pi_t / \partial s_t < 0$ is more easily satisfied. Therefore, it is more likely for the seller to cheat when he has the option to change his identity and end the game earlier.

From the above analysis, we find that neither the accumulative nor the average score mechanism provides strong incentives for the seller to consistently stay honest. With the possibility of identity change, sellers will be even more likely to cheat. One possible remedy is to increase the cost of creating a new identity. However, it will raise the cost for new entrants to the market. Thus, it may not be in the interest of online markets to implement the rule. The proposal to use free but permanent pseudonyms is a more plausible solution, as discussed in [10].

3 Exponential Smoothing

3.1 Infinite Horizon

As discussed earlier, the major problem with both accumulative and average score mechanisms is that past scores carry a large weight in determining the current reputation. Over time, the marginal benefits for the sellers to stay honest will decrease. Meanwhile, the monetary gains of cheating remain the same. When the benefits of cheating are more than the benefits of staying honest, cheating will likely occur. The goal of our design of a cooperative reputation system is to remedy this problem by controlling the impact of history. Although past transaction information is needed, we need to give more weight to the more recent activities. Such a design philosophy is, in fact, used in many rating systems, e.g., Business Week's ranking of business schools.

In our design, the overall reputation rating for the seller at time $t + 1$ is given by an exponentially smoothed function of the history scores:

$$R^s_{t+1} = (1 - \alpha) s_t + \alpha R^s_t, \quad (13)$$

where $\alpha \in (0,1)$ is the smoothing factor. A higher value of α means a higher weight on the past performance. Since R_t^s has a well-defined bound, we assume a linear bidding function from bidders and let $b_i = R_t^s v_i$. In equilibrium, the transaction price for the auction is $R_t^s v$. At time t, the seller's profit is

$$\pi_t = R_t^s v - \frac{1+s_t}{2} c_1 - \frac{1-s_t}{2} c_0.$$

Using a dynamic programming approach, the objective function for the seller is to maximize the total profits from the current period,

$$\begin{aligned} \hat{\Pi}_t(R_t^s) &\equiv \max_{s_t} \Pi_t(R_t^s, s_t) \\ &= \max_{s_t} \left(R_t^s v - \frac{1+s_t}{2} c_1 - \frac{1-s_t}{2} c_0 + \delta \cdot \hat{\Pi}_{t+1}(R_{t+1}^s) \right). \end{aligned} \tag{14}$$

The derivative of the objective function with respect to s_t is

$$\begin{aligned} \frac{\partial \Pi_t}{\partial s_t} &= -\frac{c_1 - c_0}{2} + \delta \frac{\partial \hat{\Pi}_{t+1}}{\partial R_{t+1}^s} \frac{\partial R_{t+1}^s}{\partial s_t} \\ &= -\frac{c_1 - c_0}{2} + \delta(1-\alpha) \frac{\partial \hat{\Pi}_{t+1}}{\partial R_{t+1}^s}. \end{aligned} \tag{15}$$

Similarly, we can find an iterative relation for the derivative,

$$\frac{\partial \hat{\Pi}_{t+1}}{\partial R_{t+1}^s} = v + \delta \frac{\partial \hat{\Pi}_{t+2}}{\partial R_{t+2}^s} \frac{\partial R_{t+2}^s}{\partial R_{t+1}^s} = \frac{v}{1 - \delta\alpha}. \tag{16}$$

Substituting (16) into (15), we get

$$\frac{\partial \Pi_t}{\partial s_t} = -\frac{c_1 - c_0}{2} + \frac{\delta v (1-\alpha)}{1 - \delta\alpha}. \tag{17}$$

By forcing $\partial \Pi_t / \partial s_t > 0$, we can guarantee that the seller will behave honestly. This gives us the following result.

Proposition 5. *Under exponential smoothing and an infinite horizon, if*

$$\alpha < \frac{2\delta v - (c_1 - c_0)}{2\delta v - \delta(c_1 - c_0)},$$

it is always in the seller's best interest to provide honest quality.

With this design, we can see, from (17), that the seller's decision does not depend on R_t^s and t. It is a function of the transaction-specific variables, i.e., c_1, c_0, and v, and the discount factor δ. Therefore, we can eliminate the problem that the seller has decreasing incentive to stay honest over time.

Now, we examine the comparative statistics of the solution. First, we can derive that $\partial_\delta(\partial \Pi_t / \partial s_t) > 0$ from (17). This suggests that the higher the seller values the future gain, the more likely the seller will behave honestly at the present time. If the seller does not care too much about future payoffs, he is less concerned to build a good reputation in early periods and is more likely to cheat. Second, we have $\partial(\partial \Pi_t / \partial s_t)/\partial(c_1 - c_0) < 0$. $(c_1 - c_0)$ is the gain from cheating. The comparative statics suggest that the higher the payoff of cheating, the more likely the seller will cheat.

According to Proposition 5, the solution to the problem is to lower the threshold for α. Since

$$\frac{\partial}{\partial \alpha} \left(\frac{\partial \Pi_t}{\partial s_t} \right) = -\frac{\delta v (1-\delta)}{(1-\delta\alpha)^2} < 0,$$

a lower α helps to reduce cheating behavior. The reason is that, under a small α, the present feedback carries a larger weight in computing the reputation score. In case the seller decides to behave dishonestly, his reputation score will drop quickly. The penalty in terms of a low reputation score and subsequent lower biddings from potential buyers will deter the cheating.

3.2 Finite Horizon

Next, we analyze the exponential smoothing mechanism when the timeline is finite and the seller has the option to change his identity. Let K be the number of stages left from the current period. From (14), we get

$$\begin{aligned} \frac{\partial \Pi_t}{\partial s_t} &= -\frac{c_1 - c_0}{2} + \delta(1-\alpha) \frac{\partial \hat{\Pi}_{t+1}}{\partial R_{t+1}^s}, \\ \frac{\partial \hat{\Pi}_{t+1}}{\partial R_{t+1}^s} &= v + \delta \frac{\partial \hat{\Pi}_{t+2}}{\partial R_{t+2}^s} \frac{\partial R_{t+2}^s}{\partial R_{t+1}^s} = \frac{v(1 - \delta^K \alpha^K)}{1 - \delta\alpha}. \end{aligned}$$

Therefore, we have

$$\frac{\partial \Pi_t}{\partial s_t} = -\frac{c_1 - c_0}{2} + \delta(1-\alpha) v \frac{1 - \delta^K \alpha^K}{1 - \delta\alpha}.$$

Consider the case where $K = 1$, i.e., there is only one stage to go, we have

$$\left. \frac{\partial \Pi_t}{\partial s_t} \right|_{K=1} = -\frac{c_1 - c_0}{2} + \delta(1-\alpha) v.$$

If $\alpha < 1 - (c_1 - c_0)/2\delta v$, $\partial \Pi_t / \partial s_t \big|_{K=1} > 0$. This means that the seller will provide honest quality except at the last stage. However, if $\alpha \geq 1 - (c_1 - c_0)/2\delta v$, the seller will provide inferior quality earlier. We can see that $\partial \Pi_t / \partial s_t$ is an increasing function of K. Given α and other parameters, let K^* be the threshold such that $\partial \Pi_t / \partial s_t < 0$ if $K < K^*$. If there are more than K^* stages left to go, the seller will provide honest quality.

Proposition 6. *Using exponential smoothing, if*

$$\alpha < 1 - (c_1 - c_0)/2\delta v,$$

the seller will maintain honest quality service except at the very last stage. If $\alpha \geq 1 - (c_1 - c_0)/2\delta v$, the seller will maintain honest quality if there are more than K^* stages remaining, where K^* satisfies

$$-\frac{c_1 - c_0}{2} + \delta(1-\alpha) v \frac{1 - \delta^{K^*} \alpha^{K^*}}{1 - \delta\alpha} = 0.$$

We calculated a numerical example for the model. The effect of α on seller behavior is depicted in Fig. 1. Clearly, there are trade offs between large and small α parameters. High α value means that past performance carries a higher weight in reputation rating and, therefore, reputation score

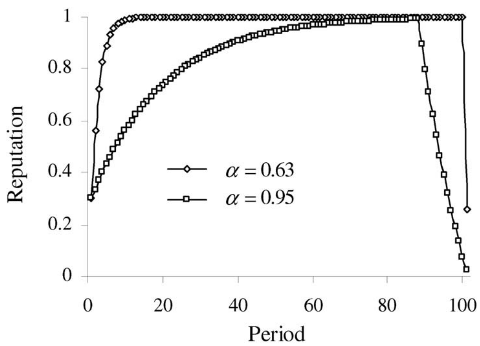

Fig. 1. Reputation scores with different α with $v = 1$, $c_1 - c_0 = 0.7$, and $\delta = 0.95$.

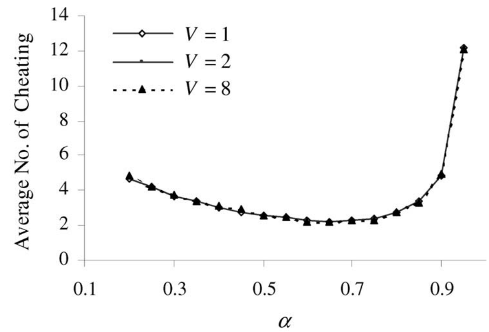

Fig. 2. Average number of cheating by α and auction item value $m \sim \mathrm{U}[0.4, 0.6]$.

tends to change slowly. It has two implications. First, when the reputation score is low, it takes a longer time for the seller to build a good reputation. Second, once the seller has a high score, it will drop slowly even if the agent cheats. On the other hand, a small α allows the reputation score to weigh heavily on the recent performance. Thus, the reputation score rises and drops quickly. It is easier for a seller to build a high reputation score. But, the seller will lose his good reputation quickly if he cheats. Those findings will be important in developing policies for the more complex situations that will be discussed in Section 4.

Despite the trade offs between large and small α parameters, our numerical results suggest that a smaller α is strictly better under a constant product value. We show the relationship between α and the system performance in Table 1. With a large α, the average number of cheating is higher and the seller cheats earlier. When α takes a value (in bold font) that is just below the threshold value as suggested in Proposition 6, the seller will not cheat except in the last period. When the auction item remains the same and α is below the threshold value, it is not worthwhile for the seller to cheat. In doing so, the seller will incur immediate losses resulting from a quick drop of the reputation score. And, the loss is unlikely to be recouped in the future because the item value is a constant.

TABLE 1
Numerical Results under Exponential Smoothing, $v = 1$, $\delta = 0.95$, and $K = 100$

$c_1 - c_0$	α	Stage Seller Starts to Cheat	Number of Cheating
0.4	**0.78**	100	1
	0.85	99	2
	0.95	95	6
0.6	**0.68**	100	1
	0.85	98	3
	0.95	91	10
0.8	**0.57**	100	1
	0.85	97	4
	0.95	84	17

4 Simulation Results

Our analysis has assumed a static value of the auction item thus far. However, the actual transactions are more complicated. From transaction to transaction, the seller usually changes the sale item and the values for the items will change as well. This provides an additional challenge in implementing the reputation mechanism. Potentially, the seller could build a high reputation by selling small-value items and decide to cheat when selling an expensive item. In this section, we assume that the value of the sale item is random, following some statistical distributions, and examine policies to implement the exponential smoothing mechanism. In addition, we explore the use of a transaction value adjusted mechanism to prevent possible malicious behavior.

4.1 Uniform Value Distribution

We first assume that the product value follows a uniform distribution, $v \sim \mathrm{U}[0, V]$, where V is the upper bound for values. The cheating margin, defined as $m = (c_1 - c_0)/v$, is assumed to follow a uniform distribution, $m \sim \mathrm{U}[m_L, m_H]$, with $0 \leq m_L \leq m_H \leq 1$. We further assume that m and v are independent, even though this assumption can be easily relaxed. The total number of stages of the game is 100.

As shown in Fig. 2, when we fix the range of the cheating margin (m), the curves representing different value distributions collapse into one. Therefore, we can safely assume that the item value follows a uniform distribution in $[0, 1]$ without the loss of generality. We find that the average number of cheating is correlated with the margin of cheating (Fig. 3). This is easy to understand because the amount of cheating increases in the margin of cheating.

From the results shown in Table 2, we can see the existence of an optimal α range that can minimize the amount of cheating. We let the margin of cheating follow a uniform distribution in $[m_L, m_H]$, i.e., $m \sim \mathrm{U}[m_L, m_H]$. We find that the optimal α ranges from 0.59 to 0.65. This result is quite different from that of the constant item auction in which there exists an optimal α threshold and a smaller α is strictly better. When the item value is dynamic, the results suggest that the optimal α cannot be too large or too small. As discussed in the last section, there are trade offs between a large and small α, which are proven to be important in this case.

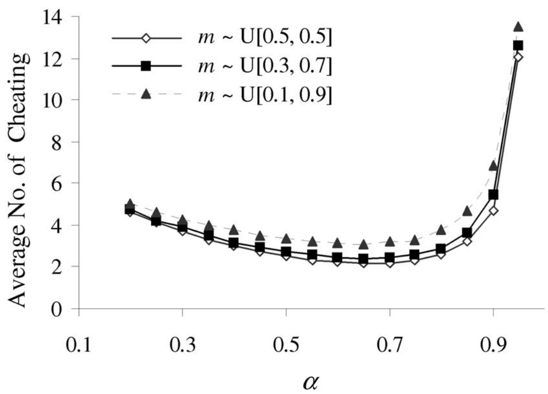

Fig. 3. Average number of cheating by α and margin of cheatings.

4.2 Lognormal Value Distribution

We next assume that the value of the sale item follows a lognormal distribution, explicitly,

$$f(v) = \frac{1}{\sqrt{2\pi}\sigma v}\exp\left(-\frac{1}{2}\frac{(\ln v - \mu)^2}{\sigma^2}\right),$$

where the mean of v is $\bar{v} = e^{\mu+\sigma^2/2}$ and the variance is $\sigma_v^2 = \bar{v}^2(e^{\sigma^2} - 1)$. In this experiment, we fix the mean $\bar{v} = 1$ and vary the coefficient of variation $(\mathrm{CV} = \sigma_v/\bar{v})$. To get a random value $\tilde{v}$, we first generate a normally distributed random number $\tilde{x}$, with mean $\mu = -\ln(1 + \mathrm{CV}^2)/2$ and variance $\sigma^2 = \ln(1 + \mathrm{CV}^2)$. Then, $\tilde{x}$ is converted to $\tilde{v}$ by $\tilde{v} = e^{\tilde{x}}$. This guarantees that $\tilde{v}$ is always nonnegative. Compared to the uniform distribution, the lognormal distribution does not set an upper bound for the product value, but, rather, has a long tail. Occasionally, the seller may sell a very expensive item. The lognormal distribution can represent the actual transactions better and presents a more challenging environment in comparison to the uniform distribution.

We also fix the margin of cheating within 0.7. As shown in Fig. 4, the simulation results suggest that the average number of cheatings goes up when the CV increases. This result is different from the case of the uniform distribution because uniform distribution's CV is a constant. The results (Fig. 4) also suggest the existence of an optimal α. We searched the optimal α value and find that the optimal α increases with the CV (Fig. 5). Overall, the optimal α value falls in a range centered on 0.5. This is consistent with the findings in the uniform value distribution.

TABLE 2
Optimal α for Different Margins (m)

Average m	m_L	m_H	Optimal α
0.4	0.4	0.4	0.59
	0.3	0.5	0.59
	0.2	0.6	0.65
0.5	0.5	0.5	0.65
	0.4	0.6	0.65
	0.3	0.7	0.65
0.6	0.6	0.6	0.65
	0.5	0.7	0.65
	0.4	0.8	0.65
0.7	0.7	0.7	0.65
	0.6	0.8	0.65
	0.5	0.9	0.65

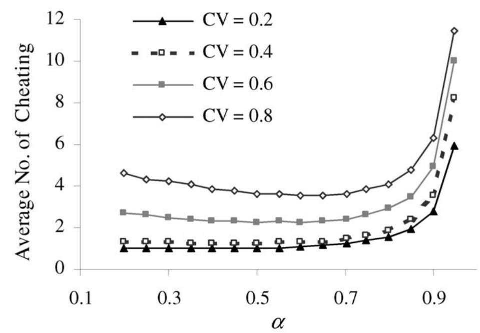

Fig. 4. Average number of cheatings by α and coefficient of variation (CV).

4.3 Optimal Policy under Perfect Information

Although there exists an optimal α that can minimize the amount of cheating for the lognormal value distribution, the rate of cheating increases quickly in CV (Fig. 4). This is a serious limitation of a single α policy. When we have a constant item value or the CV of the distribution is not too large, a single α policy provides a reasonable solution for controlling the amount of cheating, as shown in Figs. 1, 3, and 4. However, when the variance of the product value is high, e.g., in the lognormal distribution, a single α policy does not perform well. One possible solution is to have a dynamic α policy.

Let T be the total number of periods and $k\,(k = 1, 2, \ldots T)$ be the kth period before the end. With dynamic value (v), cost (c_1, c_0), and α, we can derive the following optimal policy by some modification of (15):

$$\alpha_{T-k} < 1 - \frac{\tilde{c}_1^{T-k} - \tilde{c}_0^{T-k}}{2\delta\tilde{v}_{T-k}},$$

where $\tilde{v}_{T-k} = v_{T-k+1} + \delta\alpha_{T-k+1}\tilde{v}_{T-k+1} (k = 1, 2, \ldots T)$ and $\tilde{v}_T = 0$. α_{T-k} is the optimal parameter for the unique $(T-k)th$ period.

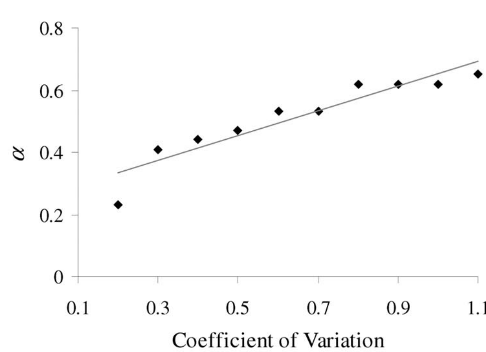

Fig. 5. Optimal α by coefficient of variation.

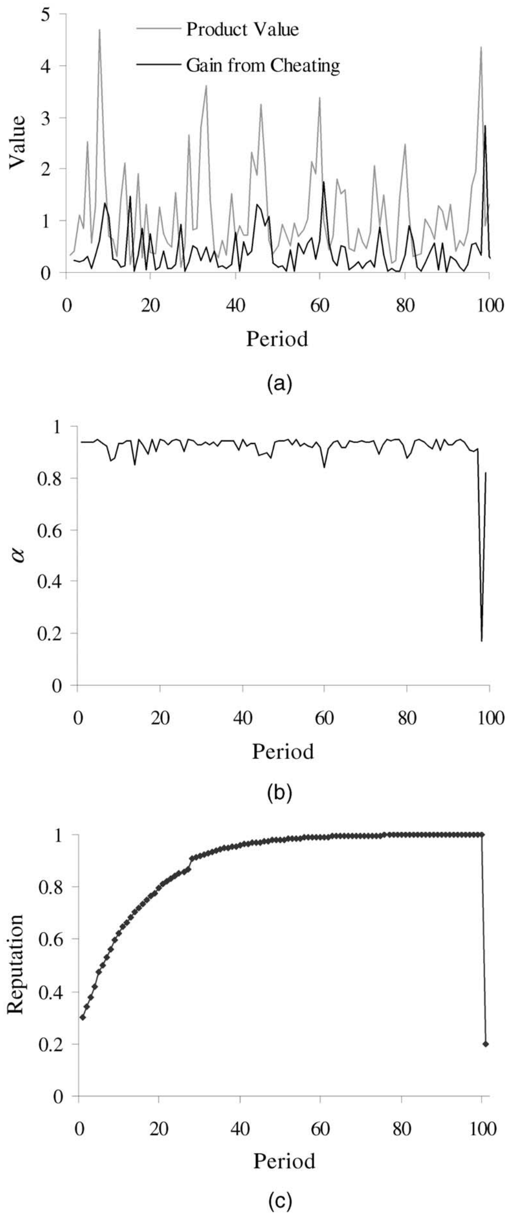

Fig. 6. Dynamic values and α.

We implement the policy by backward iterations and the results are shown in Fig. 6. By adjusting α dynamically, we are able to prevent the seller from cheating except in the last period. However, this policy suffers a serious problem—it is difficult to implement ahead of time. In order to implement the policy, it is required that we know the realizations of the value and cost distributions. Although the policy is optimal ex post, we can only find the dynamic α by assuming perfect information. The strong condition seriously limits the applicability of the mechanism.

However, by analyzing the perfect information case, we are able to derive helpful insights. First, we have observed that, in general, α is much higher than the single optimal α value that was discussed earlier. Second, α is negatively correlated with the possible gains from cheating (Fig. 6). That means when the possibility of cheating is low, α is large, which allows reputation scores to grow slowly. But,

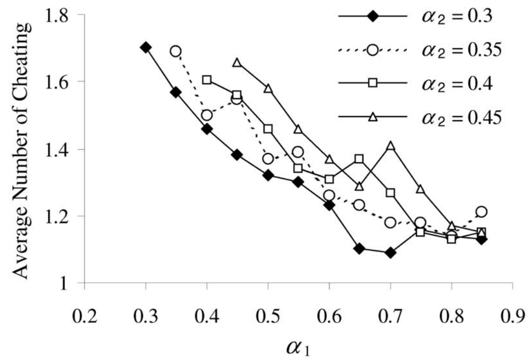

Fig. 7. Two-level α policy.

when the possibility of cheating is high, α will decrease, which makes the reputation score depend more on the current behavior. Thus, should the seller cheat, the penalty for cheating increases under a small α. We can see the asymmetry of the α parameter is effective in controlling the cheating behavior.

4.4 Two-Level α Policy

From the insights derived from the dynamic optimal α polity, we try a simple implementation without relying on the prior knowledge of the realizations of the value and cost distributions. Rather than having a dynamic α that changes in each period with product value and possible gain from cheating, we develop a two-level α policy with $0 < \alpha_2 < \alpha_1 < 1$. As long as the seller does not cheat, α_1 is applied to compute the reputation score. If the seller cheats, α_2 is applied.

This policy is intended to combine the strengths of two different levels of α parameters while limiting their weaknesses. As discussed earlier, a small α is effective in penalizing the cheating behavior with a quick plunge of the reputation score. However, the seller can regain the reputation quickly if he stays honest for a while. Therefore, a small α may induce oscillatory behavior and its effectiveness is limited. On the other hand, with a large α, the change of reputation score is slow. A seller has to work hard for a long time in order to establish a good reputation. But, once the seller has a high reputation score, the decline is also slow when he cheats. Under the design of a two-level α, we can institute an effective incentive mechanism for the seller. In general, a high α is used. Thus, it will take a long time for the seller to build a good reputation. If the seller cheats, a small α will be applied to make the punishment severe. Therefore, the seller may realize that it is not in his best interest to cheat.

We search such a combination of α_1 and α_2 in simulations and compare its performance with the single α policy. We let CV be 0.8 and $T = 20$ and conducted the simulation study. The results are shown in Fig. 7. The simulation results have several interesting findings. First, we find that the two-level α policy consistently outperforms the single α policy. Second, the optimal parameters have converged to a region with $\alpha_1 = 0.85$ and $\alpha_2 \in [0.3, 0.45]$, with the average number of cheating below 1.2. Therefore, we can safely use those parameters to implement the policy. Another advantage of this policy is that it is easy to

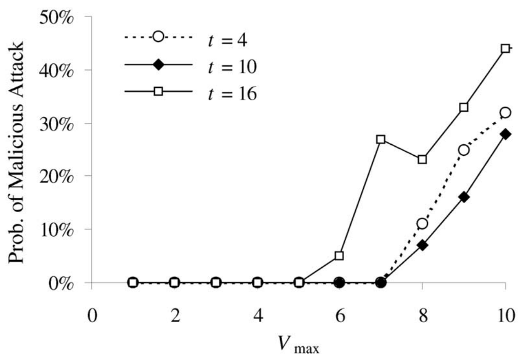

Fig. 8. Probability of malicious attack.

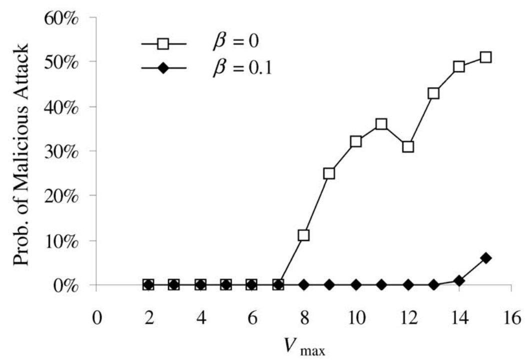

Fig. 9. Value adjusted reputation.

implement and we do not need to know the realizations of the distributions ahead of time.

4.5 Malicious Behavior Prevention

In this section, we study a possible malicious behavior of value sniping, that is, cheating on a transaction where the item value is significantly higher than the rest of the transactions.

We have set up a simulation of 20 periods of transactions. The item value is set to be 1 for all periods except the tth period, where the value is $V_{\max}(> 1)$. The margin of cheating follows a uniform distribution in $[0, 0.7]$. Fig. 8 depicts the probability of a malicious behavior that happens at the tth period. A two-level α policy is applied where $\alpha_1 = 0.8$ and $\alpha_2 = 0.3$. In general, if the item value is not too high compared to the normal, value sniping can be prevented. However, as $V_{\max}$ increases, the probability starts to increase fairly rapidly. Intuitively, a seller is more likely to cheat closer to the end of the game $(t = 16)$, as a higher reputation will not be very beneficial any more. It is also interesting to see that a seller has a higher probability of cheating early $(t = 4)$ than in the period $t = 10$. This is due to the fact that the seller can cheat and still has time to rebuild his reputation. This certainly calls for a better system that can issue a more severe penalty, for example, using a value adjusted scheme.

We have explored a simple scheme that can incorporate the information of transaction value in the reputation system. The reputation score given in (13) is modified as,

$$\bar{R}^s_{t+1} = R^s_{t+1} - \beta(V_{\max} - 1)(1 - s_t)/2.$$

If the seller cheats at the tth period (or $s_t = -1$), an additional penalty proportional to value difference is applied. Fig. 9 (where $t = 4$) shows the performances of value-adjusted scheme significantly outperforms the original one.

5 Conclusions, Limitations, and Future Research

In this paper, we first analyzed accumulative and average score mechanisms that are widely employed in the online marketplace. We find the two mechanisms have certain weaknesses. Under both mechanisms, sellers will lose incentives when their reputation scores are high enough and the transaction history is long enough. We designed a new mechanism based on exponential smoothing. This mechanism is robust in both infinite and finite horizons. In addition, we relax the constant product value assumption and allow item value to vary randomly over time. A two-level α policy is developed. The simulation studies have proven that such a policy can serve as a sustained incentive mechanism for the seller. Further, we consider the possibility of value sniping behavior and analyze the performance of the a value adjusted rating system.

This research makes several contributions. The theoretical model analyzes the existing popular reputation systems and points out the weaknesses of those systems. It will help online companies to understand the complex dynamics of the online community and develop better systems in the future. Based on the analytical model, we proposed new policies that can provide a sustained incentive for sellers.

The study has several limitations. First, our analytical reputation model assumes that the seller has a belief in the average bidder behavior. Although this assumption is supported in empirical studies [2], [12], it is a simple representation of the collective behavior of the bidders. Second, we did not include the total number of transactions as a parameter in evaluating the average score mechanism. Our purpose is to show the problem of ignoring the recency of the scores for the average mechanism. Future study can study the possible effect of the total number of transactions with the average score mechanism. Third, although the value adjusted system looks promising in preventing malicious behavior, the impact of the mechanism on regular transactions needs to be further investigated.

A reputation system is a community enforcement mechanism. In the online world, a reputation system is only effective when a dishonest seller's bad behavior is eventually punished. How well the reputation system functions depends on many factors. Our study is an attempt to analyze the existing systems and offer some improvement. There are still many important issues that need to be studied in the future. First, it is promising to deal with high variation of the sale item value by introducing value adjusted reputation scores. How to incorporate transaction value in the reputation rating requires further studies. Second, a reputation system is only effective when buyers and sellers in the marketplace actively participate in the online community and provide unbiased feedback. However, the provision of feedback is costly to the providers but benefits the whole community. It is still unclear how this voluntary contribution system works and whether it is sustainable in the long run. Third, there could be rating errors from the reviewers. The system has to be robust enough to filter out the noises.

ACKNOWLEDGEMENTS

The authors thank the associate editor, Alexander Tuzhilin, and three anonymous reviewers for their constructive comments that have greatly improved the quality of the paper. Yong Tan acknowledges with gratitude the research support by a grant from the Ford Motor Company and by the Neal and Jan Dempsey Fellowship.

REFERENCES

[1] G.A. Akerlof, "The Market for Lemons: Quality Uncertainty and the Market Mechanism," *Quarterly J. Economics,* vol. 84, pp. 488-500, 1970.
[2] S. Ba and P.A. Pavlou, "Evidence of the Effect of Trust Building Technology in Electronic Markets: Price Premiums and Buyer Behavior," *MIS Quarterly,* vol. 26, pp. 243-268, 2002.
[3] S. Ba, A.B. Whinston, and H. Zhang, "Building Trust in Online Auction Markets through an Economic Incentive Mechanism," *Decision Support System,* vol. 35, pp. 273-286, 2003.
[4] S. Braynov and T. Sandholm, "Contracting with Uncertain Level of Trust," *Computational Intelligence,* vol. 18, pp. 501-514, 2002.
[5] J.S. Coleman, *Foundations of Social Theory.* Cambridge, Mass.: The Belknap Press of Harvard Univ. Press, 1990.
[6] P. Dasgupta, P. Hammond, and E. Maskin, "The Implementation of Social Choice Rules: Some General Results on Incentive Compatibility," *Rev. Economic Studies,* vol. 46, pp. 185-216, 1979.
[7] C. Dellarocas, "Immunizing Online Reputation Reporting Systems against Unfair Ratings and Discriminatory Behavior," *Proc. Second ACM Conf. Electronic Commerce,* 2000.
[8] C. Dellarocas, M. Fan, and C. Wood, "Reciprocity, Free Riding, and Participation Decay in Online Communities," *Workshop Information Systems and Economics,* 2003.
[9] "eBay's Worst Nightmare," *Fortune Magazine,* 26 May 2003.
[10] E. Friedman and P. Resnick, "The Social Cost of Cheap Pseudonyms," *J. Economics and Management Strategy,* vol. 10, pp. 173-199, 2001.
[11] A. Greif, P. Milgrom, and B.R. Weingast, "Coordination, Commitment, and Enforcement: The Case of the Merchant Guild," *J. Political Economy,* vol. 102, pp. 745-776, 1994.
[12] D. Houser and J. Wooders, "Reputation in Auctions: Theory and Evidence from eBay," working paper, Dept. of Economics, Univ. of Arizona, 2001.
[13] J.P. Krahnen and M. Weber, "Generally Accepted Rating Principles: A Primer," *J. Banking & Finance,* vol. 25, pp. 3-23, 2001.
[14] D. Kreps, P. Milgrom, J. Roberts, and R. Wilson, "Rational Cooperation in Finitely-Repeated Prisoners' Dilemma," *J. Economic Theory,* vol. 27, pp. 245-252, 1982.
[15] D. Lucking-Reiley, D. Bryan, N. Prasad, and D. Reeves, "Pennies from eBay: The Determinants of Price in Online Auctions," working paper, Vanderbilt Univ., 1999.
[16] M. Melnik and J. Alm, "Does a Seller's Ecommerce Reputation Matter? Evidence from eBay Auctions," *J. Industrial Economics,* vol. 50, pp. 337-349, 2002.
[17] 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,* M.R. Baye, ed., Amsterdam: Elsevier Science, 2002.
[18] C. Shapiro, "Premium for High Quality Products as Returns to Reputations," *Quarterly J. Economics,* vol 98, pp. 659-679, 1983.
[19] B. Yu and M.P. Singh, "Distributed Reputation Management for Electronic Commerce," *Computational Intelligence,* vol 18, pp. 535-549, 2002.
[20] W. Vickrey, "Couterspeculation, Auctions, and Competitive Sealed Tenders," *J. Finance,* vol. 160, pp. 8-37, 1961.
[21] "Internet Fraud Increased Sharply Last Year," *Wall Street J.,* 10 Apr. 2003.
[22] C. Wood, M. Fan, and Y. Tan, "An Examination of Reputation Systems for Online Auctions," working paper, Univ. of Washington, 2002.
[23] G. Zzcharia, A. Moukas, and P. Maes, "Collaborative Reputation Mechanisms for Electronic Marketplaces," *Decision Support Systems,* vol 29, pp. 371-388, 2000.

Ming Fan received the PhD degree in information systems management from the University of Texas at Austin. He is an assistant professor at the University of Washington Business School. His research interests include economics of information technology, online community, and the use of information technology in financial services companies. He has published in *Information Systems Research*, *Communications of the ACM*, and *Computer* among other journals. He is a coauthor of the book *Electronic Commerce and the Revolution in Financial Markets*.

Yong Tan is an assistant professor and Neal and Jan Dempsey Faculty Fellow at the University of Washington Business School. His research interests include Internet caching, software engineering, and economics of information systems. He has published in *Operations Research* and *INFORMS Journal on Computing*, among other journals.

Andrew B. Whinston is the Hugh Roy Cullen Centennial Chair Professor in information systems at the McCombs School of Business, University of Texas at Austin. He is a professor in the Departments of Economics and Computer Science as well. Dr. Whinston is the director of the Center for Research in Electronic Commerce, a pioneering research facility in electronic commerce. He is the chief editor of the journals *Decision Support Systems* and *Journal of Organizational Computing and Electronic Commerce.*