Reputation, Trust, and Rebates: How Online Auction Markets Can Improve Their Feedback Mechanisms

Lingfang (Ivy) Li
School of Economics
Shanghai University of Finance and Economics
Shanghai, China, 200433
lingfangivy.li@gmail.com

Reputation systems constitute an important institution, helping sustain trust in online auction markets. However, only half of buyers leave feedback after transactions, and nearly all feedback is positive. In this paper, I propose a mechanism whereby sellers can provide rebates (not necessarily in monetary form) to buyers contingent upon buyers' provision of reports. Using a game theoretical model, I show how the mechanism can increase unbiased reporting. There exists a pooling equilibrium where both good and bad sellers choose the rebate option, even though their true types are revealed through feedback. The mechanism also induces bad sellers to improve the quality of the contract.

1. Introduction

Reputation systems that rely on feedback from traders are important institutions that help sustain trust, especially in online markets where buyers and sellers are remote and anonymous. A report by Forrester Research forecasts that online consumer auctions sales will reach US$65 billion by 2010, accounting for nearly one-fifth of all online retail

A major part of this research was carried out as one part of the author's PhD thesis at the University of California, Irvine. The paper was completed and accepted while a faculty member of the Economics Department at the University of Louisville. The author thanks the editor, coeditor, and two anonymous referees for comments on earlier drafts. The author is very grateful to Donald Saari for his encouragement and suggestions on this paper. The author thanks Yannis Bakos, Michael Baye, Luis Cabral, Jiawei Chen, Linda Cohen, Sanjeev Dewan, Per Fredriksson, Amihai Glazer, Jun Ishii, Babu Nahata, Michael McBride, Richard McLean, Stergios Skaperdas, Hal Varian, Kevin Zhu, members of the weekly discussion group at UC Irvine; and participants in the IT Economics workshop at NYU, Microeconomics Theory seminar at the University of Rutgers, EARIE, WEA, SEA, INFORMS, WISE, and Google Research. The author is grateful for the financial support from the IMBS and UC Regents. This project is also supported by Leading Academic Discipline Program , 211 Project for Shanghai University of Finance and Economics the 3rd phase. All errors are the author's.

Journal of Economics & Management Strategy, Volume 19, Number 2, Summer 2010, 303–331

sales.[1] An important reason for the success of online auction sites is that the use of an online feedback mechanism as a reputation system helps sustain trust in online markets.[2]

eBay, the largest online auction market in the United States, is the most studied online auction market in Economics and Management. Many researchers have studied the effectiveness of eBay's reputation system. As summarized in Dellarocas (2003a), the main conclusions derived from the collective body of literature show that feedback profiles seem to affect both the price and the probability of a sale. However, the precise effects are ambiguous; different studies focus on different aspects of eBay's complex feedback profile and often reach different conclusions.[3] For instance, Houser and Wooders (2006) find that a 10% rise in the number of positive ratings recorded for a seller is associated with a 0.17% increase in the price that the seller can command, whereas a 10% rise in neutral or negative ratings lowers the price received by 0.24%. Resnick et al. (2006) point out the omitted variable problem in the previous research and use controlled experiments to address this problem by holding constant the potential confounding factors in previous observational studies. They find that buyers are willing to pay 8.1% more for pairs of lots of vintage postcards from an established seller rather than from new vendors, whereas one or two negative feedback reports for new sellers do not affect buyers' willingness-to-pay.

Despite the success of online auction sites, current online reputation systems are not perfect, and online auction fraud still accounts for a significant proportion of the complaints filed with the Federal Trade Commission (FTC).[4] Asymmetric information about a seller's quality could destroy the trust and trustworthiness in online markets, limiting the market's growth.

There are several major problems in existing reputation systems, including abuse of the reputation system (or untruthful reporting), traders changing identities, low incentives for providing feedback, and bias toward positive feedback. For example, traders can manipulate their reputation scores by participating in a market for reputation on

1. Forrester Research, "US Online Auction Sales, 2005 to 2010," Report (October, 2005). Available online at http://www.forrester.com/Research/Document/Excerpt/0,7211,37848,00.html (accessed on January 28, 2006).

2. See Resnick and Zeckhauser (2002), Dellarocas (2004), and Jøsang et al. (2007). Not only auction sites but also online resellers use feedback systems to help sustain trust (e.g., Resellerratings.com and Pricegrabber.com). This paper considers a special case of auction sites.

3. See also Resnick and Zeckhauser (2002) and Bajari and Hortaçsu (2004).

4. According to Anderson (2005) online auction fraud complaints made up 41,796 out of 180,000 total complaints filed with the FTC from January 2005 to June 2005, and consistently ranked near the top of the list for all fraud complaints filed with the FTC from 2000 to June 2005.

eBay, as examined in Brown and Morgan (2006),[5] and can abuse the reputation system through bad mouthing and ballot stuffing.[6] Several papers find that buyers provided feedback on sellers about half of the time on eBay, and about 99% of the comments were positive.[7] Dellarocas and Wood (2008) find that on average, eBay buyers walk away from a transaction satisfied 78.9% of the time. One explanation for the positive bias blames missing negative feedback. Klein et al. (2005) and Dellarocas and Wood (2008) suggest that buyers are reluctant to leave negative feedback due to a fear of retaliation by strategic sellers. Strategic sellers may retaliate against buyers who leave negative feedback by leaving negative evaluations of buyers. Thus, buyers worried about their own reputation keep silent even if willing to spend time and effort on reporting their poorly performing counterparties. This explains why eBay implemented a policy to ban sellers from leaving negative feedback to buyers in May 2008.[8] However, eBay's mandatory policy made some sellers switch to other auction sites, and many sellers stopped leaving any feedback to buyers.[9] As a consequence, the mandatory policy may lower market activities as well as feedback participation rates. This paper provides a solution that works to solve the problem of missing reports and eliminates the side effects resulting from eBay's mandatory policy.

Many papers provide various mechanisms to fix the problems of low incentives for providing feedback and bias toward positive feedback[10] however, the problem of who should and would incentivize buyers to leave feedback is unsolved. This paper contributes to the scholarship on reputation mechanism by addressing the problem of who should and would provide incentives for buyers to leave feedback without requiring market makers (e.g., eBay) to provide such incentives; at the mean time, it increases feedback participation rates and solves the problems of bias toward positive feedback due to missing reports without driving out some sellers.

5. Some sellers on eBay sell low-priced, valueless items designed only to artificially enhance the seller's feedback ratings.

6. Ballot stuffing occurs when a seller colludes with other buyers to undertake fake transactions to enhance her reputation. Bad mouthing occurs when a seller is targeted by a group of buyers to deliberately lower her reputation.

7. See Resnick and Zeckhauser (2002), Cabral and Hortaçsu (2006), and Klein et al. (2005).

8. See *http://www2.ebay.com/aw/core/200801290559182.html* (accessed on September 28, 2008.)

9. See *http://money.cnn.com/2008/05/14/smbusiness/ebayfeedback.fsb/index.htm* (accessed on September 28, 2008.)

10. For example, Ba and Pavlou (2002); Dellarocas (2003b); Miller et al. (2005); Jurca and Faltings (2005, 2006, 2007), and Papaioannou and Stamoulis (2005). Please see the literature review for more detailed information on the papers.

The mechanism proposed in this paper gives sellers an *option* to provide rebates (not necessarily in monetary form) contingent on buyers leaving feedback. Using a game theoretical model, I find that the rebate incentive mechanism can increase reporting. In both a pure adverse selection model and a model with both adverse selection and moral hazard, there exists a pooling equilibrium where both good and bad sellers choose the rebate option, even though their true types are revealed through feedback. In the presence of moral hazard, the mechanism also induces bad sellers to improve the quality of the contract.

This paper makes the following contributions: First, unlike a monitoring mechanism suggested in the earlier literature (e.g., Ba and Pavlou, 2002; Dellarocas, 2003b), it provides an *option* for sellers to *signal* quality rather than *requiring* them to do so. The model thus does not increase regulation. As an "inducing" mechanism, rather than an "enforcing" mechanism, it allows both good sellers and bad sellers to coexist in the market, but it makes it possible for buyers to distinguish between them through feedback over time. Therefore, it improves market efficiency without driving out some sellers by enacting mandatory policies, such as eBay's May 2008 policy change. Second, it provides a trust mechanism that mainly involves buyers and sellers, whereas market makers only need to create the option on the web sites and need not intervene afterward. Because there is no other third-party agent involved, transaction costs are minimal, making it a *self-sustaining* mechanism. Third, based on the proposed general idea of "rebates," this paper derives conditions under which certain sellers would choose the rebates. This paper also provides suggestions on the forms of rebates, but the exact form of rebate needs to be chosen with caution by the market makers.

The remainder of the paper proceeds as follows: Section 2 places this paper in the context of the broader literature on asymmetric information and reputation mechanism models. Section 3 sets up a model to illustrate the trust problem in online trades and how online reputation systems work. Section 4 incorporates a rebate incentive mechanism into the model. Section 5 discusses possible extensions of the incentive mechanism. Section 6 concludes with implications and suggestions for the directions for future research. Proofs are included in Appendix.

2. Related Literature

There are two bodies of literature related to this paper: asymmetric information and reputation system design. Asymmetric information regarding the quality of products or sellers has a tremendous impact on market exchange, resulting in market failure.[11] Reputation and reputation mechanisms play important roles in reducing information asymmetry and inducing agents to cooperate in such settings.[12]

The two types of asymmetric information models consider adverse selection and moral hazard. Reputation mechanisms play different roles in each of these two models. Dellarocas (2003a) points out that in the case of adverse selection, the role of a reputation mechanism is to help the community learn the (initially unknown) attributes (e.g., ability, honesty, etc.) of community members; whereas in a moral hazard setting, the objective of reputation mechanisms is to promote cooperative and honest behavior among self-interested economic agents with the threat of future punishment.[13] As Cabral (2005) summarizes, typical reputation mechanism models that incorporate "reputation" (i.e., when agents believe a particular agent *to be* something) are based on Bayesian updating of beliefs and, in an adverse selection setting, possibly signaling.[14] In other reputation models, "trust" (i.e., when agents expect a particular agent *to do* something) is modeled through repeated interaction and the possibility of "punishing" off-equilibrium actions in a moral hazard setting.[15] In this paper, I first formulate a game theoretical model of adverse selection to illustrate the trust problem in online trades and then incorporate moral hazard.

This paper also incorporates the literature on reputation system design. In an online market, a reputation system is the main institution to induce traders to behave cooperatively. Because the online reputation system controls the content and format of the aggregated information it publishes, it is important to design an incentive mechanism that will elicit unbiased feedback.[16]

Following Resnick and Zeckhauser (2002), who demonstrated the problems of low feedback rates and potential reporting biases, several papers have put forward mechanisms to solve these problems. One solution employs monitoring mechanisms. Ba et al. (2002) suggest a

11. See Akerlof (1970), Klein and Leffler (1981), Shapiro (1982, 1983), and Kauffman and Wood (2000).

12. See Fudenberg and Levine (1992), Kreps and Wilson (1982), Celentani and Pesendorfer (1996), Battigalli and Watson (1997), Levine and Martinelli (1998), and Kandori (1992).

13. For example, in the form of lower bids following the posting of a negative rating on a trader's reputation profile.

14. See Klein and Leffler (1981) and Shapiro (1983).

15. See Kreps and Wilson (1982), Milgrom and Roberts (1986), and Diamond (1989).

16. This paper focuses on the bias caused by missing (negative) reports. I do not consider "mis-reporting" (i.e., untruthful reporting) in the model, but in the discussion section, I suggest possible rebate forms that potentially may solve the mis-reporting problem as well.

Trusted Third Party (TTP) mechanism to issue certificates to sellers and buyers. Their mechanism can induce cooperative behavior if both buyers and sellers obtain the verifications from the TTP. Dellarocas (2003b) proposes charging a listing fee contingent on a seller's announced expected quality and rewarding the seller based on both his announced quality and the rating posted for that seller by the winning bidder for that listing. Both mechanisms discourage sellers from lying about the true quality of the product, thus improving the efficiency of the market. The proposed rebate mechanism in this paper has several novel features. First, the rebate mechanism is a voluntary signaling mechanism rather than a monitoring mechanism. Second, because it involves buyers and sellers without introducing other third parties,[17] the transaction costs associated with the incentive mechanism are minimal. Third, instead of imposing costs, the mechanism provides buyers incentives to leave feedback.

Another kind of mechanism design focuses on eliciting truthful feedback through peer-provided feedback. Miller et al. (2005) proposes a peer-prediction method that compares the likelihood assigned to a reference rater's possible ratings to the reference rater's actual rating. In Computer Science literature, Jurca and Faltings (2005, 2006, 2007) and Papaioannou and Stamoulis (2005) also propose well-designed systems of rewards and punishments that induce both sellers and buyers to report truthfully. The earlier mechanism's feedback elicits truthful reputation history about the online traders, thus improving the market efficiency. In all of these proposed designs, either the market makers (e.g., eBay) provide the awards, such as monetary incentives or points to the raters, or the market makers collect entry fees from raters in advance to make them voluntarily transfer points to the reference raters.[18]

This paper differs from the truth-eliciting literature in the following ways: First, this paper focuses on fixing the *positive bias* caused by the *missing negative feedback* as observed in Dellarocas and Wood (2008) and Klein et al. (2005), rather than *untruthful feedback*. It seems

17. The rebate mechanism also involves the market maker providing the option, but no other third parties get involved.

18. On page 202, Jurca and Faltings (2007) state that "the reputation system pays buyers for the submitted reports." On page 1361, Miller et al. (2005) state that "The center makes transfers to each rater, awarding or taking away points based on the raters' messages," and "though all transactions actually occur between raters and the center, this creates the effect of having the raters settle the transfers among themselves," "one way to assure ex post voluntary participation is to collect bonds or entry fees in advance, and use the collected funds to ensure that all transfers are positive." They continue on page 1364, saying"alternatively, it may be sufficient to threaten to exclude a rater from future participation in the system if she is unwilling to act as a rater or settle her account after a negative outcome" (untruthful reporting that lower the points in her account in the scoring system.)

to be unlikely that the positive bias is caused by buyers untruthfully reporting "positive," if they are not satisfied with the transactions or because of sellers' friends' ballot stuffing all the time. The auction markets also have rules that hold buyers accountable.[19] Second, this paper focuses on the seller's decision on whether to choose the rebate option as a signal and who would choose it, whereas the truth-eliciting literature focuses on how much incentive or punishment should be given to buyers to make them tell the truth. This paper compliments the truth-eliciting literature by solving the problem of who would provide the incentives. Third, if as several papers suggest, missing negative feedback is due to fear of retaliation,[20] then merely giving award points or monetary incentives may not overcome buyers' fear of speaking out and solve the positive bias problem. In Section 5.2, this paper proposes an automatic feedback option as a form of rebate to get rid of buyers' concern about or fear of retaliation. Fourth, this paper suggests giving an *option* to *sellers* to provide incentives for buyers' reports. It does *not* require the *market makers* to provide incentives. Because there are millions of transactions in online auctions, the cost of providing such incentives might be a concern for the market makers. Also, it does *not* require *buyers* to pay any cost, such as entry fees, or threaten them to participate in the rating system, rather it provides incentives for buyers to report. What is more, it does *not* require *all sellers* to choose the option. This paper explains the conditions to determine who would self-select to choose the rebate option. Because it is an option for sellers to provide incentives to encourage buyers' report, it does not impose mandatory additional cost on either buyers to report.[21]

The third kind of mechanism accounts for the missing reports through a computational mechanism. For example, Dellarocas and Wood (2008) provide a sophisticated computational mechanism to remedy distortions introduced by reporting bias. Dellarocas and Wood (2008) explain"feedback mechanism designers and users have two *nonrival* paths of action. The first path is to explore to what extent changes in the design of a feedback mechanism can reduce reporting bias. The second path is to develop methodologies that can help mechanism users make better inferences from feedback provided by today's imperfect mechanisms." Their paper deals with the second path, whereas this paper deal with the first. Like our model, their method assumes that the reporting bias is due to missing reports, and there is no

19. For example, http://www2.ebay.com/aw/core/200801290559182.html.
20. See Dellarocas and Wood (2008), Klein et al. (2005), and Li (2008).
21. For other issues on untruthful reporting, Bhattacharjee and Goel (2005), and Dellarocas (2004) suggest solutions to the problems such as bad mouthing and ballot stuffing, and Friedman and Resnick (2001) and Dellarocas (2005) suggest imposing a cost of changing IDs to solve the ease of changing online identities problems.

systematic strategic misreporting. Their mechanism takes the missing reports into consideration to "see through" the distortion introduced by reporting bias. Their mechanism *indirectly* infers feedback information from missing reports by computational mechanism, whereas this paper suggests an incentive mechanism to induce buyers to leave feedback *directly*, thus reducing missing reports. Like our model, their mechanism requires the market maker to make some changes.[22] However, unlike ours, their mechanism is much more computationally involved for nonresearchers. For example, they use maximum likelihood estimation (MLE) method to estimate the model, and the MLE method may generate biased estimates when samples are small. What is more, if the honest reporting assumption is relaxed, then their model may not produce unbiased estimate. This paper does not solve dishonest reporting problems either, but the mechanism proposed in this paper could potentially combine with truth-eliciting mechanisms to reduce dishonest reporting.

The mechanism proposed in this paper bridges the gap between the issues of incentives and who would provide those incentives, thus complimenting the existing mechanisms. The following section illustrates the asymmetric information problem in online auction markets.

3. Model of Adverse Selection and Reputation System in Online Auction Markets

In this section, I formulate a model to illustrate the trust problem of reputation systems, using eBay's reputation system as an example, with adverse selection for geographically dispersed anonymous participants in online auctions. Because most online transactions require receipt for payment before the product is sent, in the model I only consider the case where sellers have an incentive to commit fraud.[23]

3.1 Model Setup

To capture the essence of the online reputation system, the model proceeds as the following: Suppose there are M sellers and N buyers

22. In their case, the market maker needs to provide the information on the total number of transactions for each seller. In our case, the market maker needs to provide the rebate options to sellers.

23. Most complaints filed with the FTC as Internet auction fraud report problems with sellers who "fail to send the merchandize," "send something of lesser value than advertised," "fail to deliver in a timely manner," or "fail to disclose all relevant information about a product or terms of the sale." For more information, please see http://www.ftc.gov/bcp/conline/pubs/online/auctions.htm (accessed on January 20, 2006).

in the entire market where $N \gg M$, and the sellers live infinitely. This paper focuses on one auction listing where a seller (s) lists the same good (g) in each period. There are two classifications for the transaction: high-quality level (Q_H) and low-quality level (Q_L). High-quality means the good is received by the buyer, the quality of the good is the same as the seller promised, and the good is shipped on time. A transaction is of low quality if it fails any of these criteria. For each period, there are K buyers, randomly drawn with replacement from N potential buyers. The valuation of the good by the K buyers is uniformly distributed from 0 to 1, that is, $0 < V_b(1) < V_b(2) < \cdots < V_b(K) \leq 1$. The buyer K wins the bidding, and the price is settled at $P \approx P(V_b(K))$.[24] Without loss of generality, I assume that a good is worth 1 to the winning bidder if it is Q_H, and the good is worth 0 if it is Q_L. In this setting, the price of an item is determined by the reputation of the seller, and both the seller and the buyer know how the price is determined.[25] After the bidding process, there is only one seller (s) and one buyer (b) in the transaction.

Suppose there are two types of sellers, a good type (G) and a bad type (B). The commonly known proportion of good sellers is μ_0, and the proportion of bad sellers is $1 - \mu_0$. The probability that a high-quality transaction is provided by a good seller is α, and it is β for a bad seller $(0 \leq \beta < \alpha \leq 1)$.[26] In this adverse selection model, the transaction outcomes depend on sellers' types, so the sellers do not control the transaction outcomes. For simplicity, the production cost for both types of sellers is assumed to be 0, because the transaction outcome is not affected by their effort in this model.[27]

Several assumptions are made in the model. First, both buyers and sellers are risk neutral. Second, buyers meet the same seller only

24. I use eBay as an example of the online auction market, and eBay uses the Vickrey auction method, that is, the winning bidder pays the second highest bid, so $P = P(V_b(K - 1)) < P(V_b(K))$. I assume that there are many bidders for each auction, the winning bid is ξ higher than the second highest bid, and the winning bid equals the willingness to pay for the winning bidder, so I use the highest bid to approximate the willingness to pay of the winning bidder, that is, $P \approx P(V_b(K))$. It is also possible to use the auction model in Cabral and Hortaçsu (2006) to show that the winning bid is an increasing function of the buyer's willingness to pay.

25. The model is set up to appeal to the intuition coming from eBay. However, the model is exactly the same for the listed price markets, such as pricegrabber.com and resellerratings.com. For example, a seller sets a fixed price for the item at the buyers' highest willingness to pay in that market.

26. For example, the good sellers are very careful and honest, so they advertise the product according to its true condition, pack the goods carefully, and ship it on time, whereas the bad type sellers are lazy and dishonest.

27. The general result does not change as long as the cost for both types is the same. In the mixed model with moral hazard, I introduce different production costs for the two types. In Section 5.1, I consider cases where sellers can choose an effort level that will affect the transaction outcomes, so that both adverse selection and moral hazard are combined.

once.[28] Third, if there is no net reporting cost, buyers will choose to report honestly, that is, (GR) for Q_H and (BR) for Q_L. This assumption follows from the second assumption because a buyer who only buys from a seller once has no incentive to report dishonestly.[29] Fourth, the net reporting cost is C for all buyers.[30]

Below is the formal description of how the game proceeds. Each period t consists of a sequence of moves in the following order:

1. Nature chooses the seller's type $\theta \in \{\theta_G, \theta_B\}$. The seller's type is chosen in the first period, and it persists for the rest of the game.
2. The buyer chooses a bid at the expected value of the transaction, $P \geq 0$.
3. The seller chooses to accept or reject P based on his reservation price. If he rejects, the game ends. If he accepts, then the game continues to the next step. For simplicity, I assume the good seller's reservation price is $V_S^G \geq 0$, and the bad seller's reservation price V_S^B is 0.[31]
4. Nature chooses the quality of the transaction that buyers get from different types of sellers, Q_H or Q_L. Transaction quality is a new draw in every period.
   $q(\theta)$= probability of providing Q_H.
   $q(\theta_G) = \alpha; q(\theta_B) = \beta$
   $0 \leq \beta < \alpha \leq 1$
5. After the transaction, the buyer chooses to review the transaction {NR, GR, BR}. Buyers can choose to give a good report/positive feedback (GR), a bad report/negative feedback (BR), or no report (NR).
6. Payoffs received for period t are
   $U_s(\text{Accept}) = P - 0 = P; U_s(\text{Reject}) = 0;$

28. According to Resnick and Zeckhauser (2002), 89% of all buyer-seller pairs conducted just one transaction during the five-month period covered by the data set.

29. For example, a one-time buyer has no incentive to maliciously report negative to lower the seller's price in the future. In Section 6, I discuss the possible ways to apply the proposed mechanism with dishonest report assumption.

30. C is net reporting cost that equals total reporting cost minus the reporting benefit. For example, the reporting cost may be the time and effort to leave feedback, and the benefit may be the self-satisfaction of altruism. The net reporting cost is proportional to the price P. It is possible to assume the net reporting cost is uniformly distributed between $-C$ (gain more benefit through leaving a report than reporting cost) and C. To induce everyone to report (or a certain representable proportion of traders to report), it is still needed to overcome the highest net reporting cost C. Here, we assume everyone's reporting cost is C for simplicity reasons.

31. For the listed price case, in the second move the seller sets a price at the highest buyer's expected value of the transaction, if it is higher than his reservation price; otherwise, the seller sets the price at his reservation price. In the third move, the buyer chooses to buy or not to buy.

$$U_b(\mathrm{P}, \mathrm{NR}, Q_H) = 1 - P; U_b(\mathrm{P}, \mathrm{NR}, Q_L) = -P;$$
$$U_b(\mathrm{P}, \mathrm{GR}, Q_H) = 1 - P - C; U_b(\mathrm{P}, \mathrm{BR}, Q_L) = -P - C.$$

The payoff for the seller is the price he receives: 1 on the item if he accepts the bid and 0 if he rejects. If the buyer chooses no report, then his payoff is the difference between the value of the transaction (1 for high quality and 0 for low quality) and the price he pays; if he chooses to report, then the reporting cost is subtracted from his payoff.

3.2 Equilibrium

Under the reputation system, in each period the buyers who are new to the seller but not necessarily new to the market observe the reputation history of the seller. Equilibrium outcomes vary with the reporting cost. To see the impact of reporting cost on the reputation system, I consider three variations of the reporting costs: no reporting cost, symmetric reporting cost, and asymmetric reporting cost.

3.2.1 Case 1: No Reporting Cost (C = 0)

A buyer's willingness to pay in period $t + 1$ is $P_{t+1} = \mu_t\alpha + (1 - \mu_t)\beta$, where μ_t is the belief of meeting a good seller in period t. For instance, in period $t = 1$, μ_0 represents the prior belief, and the buyer's willingness to pay is $P_1 = \mu_0\alpha + (1 - \mu_0)\beta$. If the buyer receives a Q_H product and reports GR in period $t = 1$, the buyer in period $t = 2$ observes the reports and updates his beliefs on the seller's type by Bayes' rule, so the prior in period $t = 2$ is

$$\begin{aligned}
\mu_1 &= \Pr(\theta_G \mid \mathrm{GR}) \\
&= \frac{\Pr(\mathrm{GR} \mid \theta_G)\Pr(\theta_G)}{\Pr(\mathrm{GR} \mid \theta_G)\Pr(\theta_G) + \Pr(\mathrm{GR} \mid \theta_B)P(\theta_B)} \\
&= \frac{\mu_0\alpha}{\mu_0\alpha + (1 - \mu)\beta}.
\end{aligned}$$

If the buyer receives a Q_L product and reports BR in period $t = 2$, the buyer in period $t = 3$ observes the previous reports and updates his beliefs according to Bayes' rule, so the prior in the period $t = 3$ is

$$\begin{aligned}
\mu_2 &= \Pr(\theta_G \mid \mathrm{GR}, \mathrm{BR}) \\
&= \frac{\Pr(\mathrm{BR} \mid \theta_G)\Pr(\mathrm{GR} \mid \theta_G)\Pr(\theta_G)}{\Pr(\mathrm{BR} \mid \theta_G)\Pr(\mathrm{GR} \mid \theta_G)\Pr(\theta_G) + \Pr(\mathrm{BR} \mid \theta_B)\Pr(\theta_B)} \\
&= \frac{\mu_0\alpha(1 - \alpha)}{\mu_0\alpha(1 - \alpha) + (1 - \mu_0)\beta(1 - \beta)},
\end{aligned}$$

where $\Pr(\theta_G) = \mu_0$, $\Pr(\theta_B) = 1 - \mu_0$, $\Pr(\text{GR} \mid \theta_G) = \alpha$, $\Pr(\text{BR} \mid \theta_G) = 1 - \alpha$, $\Pr(\text{GR} \mid \theta_B) = \beta$, $\Pr(\text{BR} \mid \theta_B) = 1 - \beta$, and the reports are independent of each other, so $\Pr(\text{GR, BR} \mid \theta) = \Pr(\text{GR} \mid \theta)\Pr(\text{BG} \mid \theta)$.

Let t_{GR} represent the number of positive feedback and t_{BR} be the number of negative feedback to the seller prior to the period t. So for $t = t_{\text{GR}} + t_{\text{BR}} + 1$, the buyer's belief of meeting a good seller in period t is

$$\mu_{t-1} = \mu_{t_{\text{GR}}+t_{\text{BR}}} = \frac{\mu_0 \alpha^{t_{\text{GR}}}(1-\alpha)^{t_{\text{BR}}}}{\mu_0 \alpha^{t_{\text{GR}}}(1-\alpha)^{t_{\text{BR}}} + (1-\mu_0)\beta^{t_{\text{GR}}}(1-\beta)^{t_{\text{BR}}}}.$$

The seller's payoff in period t is $U_s = P = P(V_b(K)) = \mu_t\alpha + (1-\mu_t)\beta$, and the buyer's payoff in period t is $U_b(\text{P, NR/GR/BR}) = 1 - P = 1 - (\mu_t\alpha + (1-\mu_t)\beta)$ if he received high quality transaction, and it is $U_b(\text{P, NR/GR/BR}) = -P = -(\mu_t\alpha + (1-\mu_t)\beta)$ if he received low-quality transaction. As the game repeats, μ converges to 1 for good sellers and to 0 for bad sellers; buyer's willingness to pay converges to α for good type sellers and to β for bad type sellers.

No reporting cost is the ideal situation, but in reality, there always exists some kind of cost associated with reporting. I examine the cases where the costs of providing positive and negative feedback are the same and different in the following two subsections.

3.2.2 Case 2: Symmetric Reporting Cost ($C_{\text{GR}} = C_{\text{BR}} > 0$)

In this case, there are the same reporting costs on both positive (GR) and negative feedback (BR). The reporting costs may include time or effort spent on writing reports or the opportunity cost during that time. For simplicity, all reporting costs are denoted in dollars, and the highest reporting costs for the buyers is C.

No-reporting (NR) dominates reporting (GR and BR) for buyers. When there is no report about the seller's previous history and reputation, a buyer cannot update her beliefs, so the buyer's willingness to pay is $P = P_1 = \mu_0\alpha + (1-\mu_0)\beta$ for every period t. In equilibrium, the buyers' strategy is to choose to bid at P_1 and not report, that is, $(P_i = P_1, \text{NR})$, and the seller's strategy is to accept the bid if P_i is equal or greater than the seller's reservation price V_s. If assuming the reservation price is 0 for both types of sellers, they will always accept the bid. However, a good seller gets less than what she could get in the case of no reporting cost, and a bad seller gets more than what she could get in the case of no reporting cost in the long run. There is no welfare loss in this case but a wealth transfer from good sellers to bad sellers.

If a good seller's reservation price is higher than P_1, then only a bad seller accepts the bid. Good sellers will be driven out of the market, and buyers' willingness to pay will drop to $P_i = \beta$. As a result, those buyers

TABLE I.
RANKINGS OF REPUTATION PROFILES

Seller	Feedback +1	Feedback 0	Feedback −1	Total Rating
A	18	18	27	−9
B	10	30	9	1
C	21	1	13	8

who want cheap things and do not care about low-quality transactions will stay in the market with reporting cost, and those who care about the quality of transactions will go to another market where there are many good sellers.

It is in eBay's (or other online auction market makers') best interest to identify good sellers from bad. If the population of the seller is fixed at M and eBay's revenue is an increasing function of sale value, then eBay's revenue is proportional to $\mu_0 M\alpha + (1 - \mu_0)M\beta$ when there is no reporting cost. If there is a reporting cost and the good seller's reservation price is higher than P_1, then the good sellers will not sell in this market. Only bad sellers remain, so the buyers' willingness to pay will be β. If μ_0 of the M sellers are good sellers and eBay's revenue is an increasing function of sale value, then eBay's revenue is proportional to $(1 - \mu_0)M\beta$, which is less than the earlier case. In order to retain the good sellers whose reservation price is higher than P_1 in the market, we must provide means to increase buyers' willingness to pay of the good sellers. Therefore, the market needs reliable reputation systems to help to identify good sellers.

3.2.3 ASYMMETRIC REPORTING COST ($C_{GR} \neq C_{BR} > 0$)

This section examines what happens if buyers incur different costs for submitting good and bad reports. Suppose there are three sellers in the market; Table 1 shows their ratings profiles on eBay.

Seller Ann has 18 positive, 18 neutral, and 27 negative ratings, so her overall rating is −9; seller Bob has 10 positive, 30 neutral, and 9 negative ratings, so his overall rating is 1 (On eBay, the overall rating is the sum of all the positive, neutral, and positive ratings); and seller Cindy has 21 positive, 1 neutral, and 13 negative ratings, so her overall rating is 8. Thus, we see seller Cindy has the highest overall rating, followed by Bob, and then Ann (i.e., $Cindy \succ Bob \succ Ann$). However, if there are reporting costs only on good reports, GR, then we will not observe GRs, and the overall ranking will reflect that Bob is a better seller than Cindy and Ann, respectively (i.e., $Bob \succ Cindy \succ Ann$). If there are reporting costs only on bad reports, BR, then the overall ranking will be $Cindy \succ Ann \succ Bob$.[32]

This example shows what can go wrong if we do not have complete information about a trader's reputation. Using Bayesian updating, if buyers know that there are reporting costs, regardless of whether the cost is for good feedback (GR) or negative feedback (BR), they cannot correctly update their beliefs from rating profiles, so their willingness to pay is still P_1. Therefore, a mechanism is necessary that can change the last two cases to the no reporting cost case.

4. Improvement of Current Reputation Systems in Online Auction Markets

In this section, I propose an improvement on the current reputation systems in online auction markets. Namely, I propose that sellers be allowed to reward buyers who report the quality of transactions.

If there are reporting costs, previous theoretical results predict that no buyer will want to report. In reality, many buyers do report. It might be because of altruism, community recognition, social norms, emotional expression, or that different buyers have different reporting costs (Bolton et al., 2003; Gazzale, 2005; Xiao and Houser, 2005; Resnick et al., 2006). However, as long as some of the buyers bear some reporting costs, information about the sellers' reputation profiles are incomplete. Incomplete information could lead to poor bidding decisions; thus, complete or at least unbiased information about the sellers' past history is crucial. This paper suggests providing incentives for the buyer to report and making all the information available to everyone.

How might this be achieved? One way is to eliminate the reporting costs for the buyers. Will eBay want to compensate buyers for their reporting costs? It would seem prohibitively expensive, because providing even US$ 0.01 per transaction would be a huge cost for eBay due to the millions of daily transactions. What about sellers providing incentives to buyers? Will sellers want to compensate the buyers' reporting costs? What type of sellers would be willing to compensate

32. In fact, when we can observe all the information, the overall rating is consistent with the ranking by using a Borda Count voting rule. When we cannot observe GRs, the ranking is the same as using the antiplurality voting rule. When we cannot observe BRs, the ranking is the same as using the plurality voting rule. Much research has shown that the Borda Count has the fewest problems among those voting rules (Saari, 1999; Saari, 2001; Saari and Sieberg, 2001). The best case is when we can observe all the reports and sum them up. Because traders can only observe the existing reports, they do not know the total number of transactions.

these costs to the buyers? To answer these questions, we make the following analysis of the equilibrium.

4.1 Model Setup

Suppose all sellers can choose a rebate, r, which is greater than or equal to C. The game played in every period is similar to the game in Section 3.1, except that a seller can choose to rebate before buyers bid. The formal game is described as the following:

1. Nature chooses the seller's type $\theta \in \{\theta_G, \theta_B\}$.
2. The seller chooses to rebate r or not; where $r > C$, the seller's actions set is {Rebate(R_S), NoRebate(NR_S)}.

3.–6. Same as Steps 2–5 in the model of Session 3.1.

7. These are payoffs received for period t.
   U_s(Accept, R_S, if buyer chooses GR or BR) $= P - 0 - r = P - r$;
   U_s(Accept, R_S, if buyer chooses NR) $= P$; U_s(Accept, NR_S) $=$
   $P - 0 = P$;
   U_s(Reject, R_S) $= 0$; U_s(Reject, NR_S) $= 0$;
   $U_b(P, NR, if\ Q_H) = 1 - P; U_b(P, NR, if\ Q_L) = -P$;
   U_b(P, GR, $if\ Q_H$ and seller chooses R_S) $= 1 - P - C + r$;
   U_b(P, GR, $if\ Q_H$ and seller chooses NR_S) $= 1 - P - C$;
   U_b(P, BR, $if\ Q_L$ and seller chooses R_S) $= -P - C + r$;
   U_b(P, BR, $if\ Q_L$ and seller chooses NR_S) $= -P - C$;

In this model, sellers need to consider giving rebates as part of their strategy. If a seller decides to provide the rebate, then his payoff is less by r than the model in Section 3.1. For a buyer, his payoff is r higher than the case in Section 3.1 if the seller chooses the rebate. Otherwise, his payoff is the same as in the previous model.

To find possible equilibria, I use the guess and verify method to look for the Perfect Bayesian Equilibrium (PBE). As we have seen in Section 3, good sellers are more likely to sell the same product at a higher price if there are full reports, whereas bad sellers are more likely to get lower prices if there are full reports available to consumers. At first glance, it seems that good sellers would be more inclined to provide a rebate than the bad sellers. Is it possible to have such a separating equilibrium?

4.2 Separating Equilibrium

Proposition 4.1. *No separating equilibrium exists where good (bad) type sellers choose to provide feedback rebates and bad (good) type sellers choose not to do so.*

Proof. See Appendix.

The intuition of this proposition is as follows: If good sellers choose the rebate option but the bad sellers do not, then buyers can immediately identify sellers' types by observing who provides rebates. Therefore, the intuition mentioned earlier is not an equilibrium. Using similar logic, we can reject the separating equilibrium where only bad sellers provide rebates.

Having examined the separating equilibrium, the natural next step is to examine pooling equilibrium.

4.3 Pooling Equilibrium

Proposition 4.2. *If the reporting cost C is not greater than* $(1-\mu_0)(\alpha-\beta)$*, which is the expected loss of treating a bad seller as a good one, there exists an equilibrium where good sellers always provide rebates and bad sellers mimic good sellers by providing rebates until their bad types are identified through the reports and are no longer profitable to mimic good sellers, given the off-equilibrium path belief that anyone who does not choose to rebate must be a bad seller. If the reporting cost C is greater than* $(1-\mu_0)(\alpha-\beta)$*, then both types of sellers would choose not to rebate,* NR_S.

Proof. See Appendix.

The intuition behind this proposition is as follows: Once a seller chooses a rebate, it can overcome the buyer's reporting cost. Because the price is decided by an auction, then the buyer will incorporate the rebate into his bidding price to win the bid. Therefore, he will bid higher by r. As shown in the proof provided in Appendix, the final payoff of a buyer is the same as in the benchmark model where there is no reporting cost (see in Section 3.2.1), and the seller's payoff is only less by C than his payoff in the benchmark model. In other words, as long as the rebate can cover the reporting cost, then the buyers' reporting cost will transfer to the sellers who choose the rebate option. The only cost that a seller will bear is C (not r), if he chooses the rebate option.

Because a buyer is willing to pay α to a good seller and β to a bad one, if he mistakes a bad seller for a good one, his expected loss is $(1-\mu_0)(\alpha-\beta)$, because the probability of meeting a bad seller is $1-\mu_0$. This is the highest possible loss that a buyer will have without any reputation information about a seller, so it is the highest premium for a buyer's willingness to pay a good seller. If the reporting cost C is higher than this, then it is not profitable for a seller to provide rebates to cover the reporting cost. If the reporting cost is lower than this, then good sellers will provide rebates, and bad sellers will mimic good sellers until it is no longer profitable (i.e., $P - C < \beta$). After that point, a bad seller will not provide rebates but sell the product at β or change his ID to start as a new seller again, if possible. Friedman and Resnick (2001) and Dellarocas (2005) suggest ways to prevent changing IDs, so in this model, bad sellers will sell the product at β when it is not profitable to provide a rebate. In fact, this option mechanism is better than requiring all sellers to provide a rebate because that would drive bad sellers out of the market, causing the total revenue of the market maker to decrease. By using the option mechanism, bad sellers will still stay in the market but only sell the product at the low price, and buyers who want to buy cheap products will still buy from them. Therefore, the market is efficient in the sense that buyers can learn sellers' reputations and bid accordingly, but both bad and good sellers will coexist in the market and help the market grow.

Another pooling equilibrium examined is when both types of sellers do not provide rebates, and this is not sustainable based on intuition criteria to refine the PBE.[33]

Proposition 4.2 says that if the reporting cost is not too high relative to the difference between what buyers are willing to pay to the good and bad sellers, then both types of sellers will provide the rebate option. If the reporting cost is too high, then no seller wants to provide the rebate. For example, if a good seller sells a book that is worth US$ 10 in a bookstore, then the difference of willingness to pay to a good seller and to a bad seller may be just US$0.50. If the reporting cost is higher than US$ 0.50, then even the good seller does not want to give a rebate contingent upon receiving the feedback. However, if the product is a Dell computer worth US$ 1300 on Dell's web site, the difference in willingness to pay between the different types of sellers could be US$200. The reporting cost for most people is certainly less than this amount, so the good as well as the bad sellers want to provide the rebate. Our model assumes sellers list the same product in every period, so the decision of providing rebates depends on the relative cost of reporting.

5. Extensions

This section discusses possible extensions of the rebate incentive mechanism. First, I introduce moral hazard in the reputation model and show that the rebate incentive mechanism can induce bad sellers to

33. See Cho and Kreps (1987).

behave cooperatively. Second, I discuss the possible forms of rebates, how to combine rebates with other mechanisms to solve the untruthful reporting problem, and what the "real" cost is for market makers to implement the suggested mechanism.

5.1 Model of Adverse Selection and Moral Hazard

In the model of adverse selection, the transaction outcomes are chosen by nature and not by the sellers. In reality, we often see that the sellers' actions have an impact on the transaction outcomes. If the sellers put more effort into packaging and shipping the product, the transaction outcomes are more likely to be good. To model the sellers' efforts in the model, a model with both adverse selection and moral hazard is explored.

Using a set up similar to the model in Section 3.1, if both good and bad types of sellers put forth effort ($e = 1$), they will provide high-quality products (Q_H) with probability 1.[34] If they do not put forth an effort ($e = 0$), then they will provide a low-quality product (Q_L) with probability 1. Assume that good sellers' costs of making an effort is 0, $C_{\theta_G}(e = 1) = e(0) = 0$, and bad sellers' costs of making an effort is $C_{\theta_G}(e = 1) = e(1) > 0$. To simplify, good sellers are assumed always to make an effort because it costs nothing to them.

To make the game closer to reality, it is only played in T periods instead of infinitely. The game played in every period is similar to the game in Section 4.1, adding step 4 for a seller to choose the effort after he accepts the price. The game is described as the following:

1.–3. Same as steps 1–3 in the model of Session 3.1.
4. The seller chooses to put forth an effort or not, $e = 1$ or $e = 0$.
5. Same as step 5 in the model of Session 3.1.
6. The following are the payoffs received for period t.
$U_s(\theta_G, e = 1) = P; U_s(\theta_G, e = 0) = P;$
$U_s(\theta_B, e = 1) = P - e(1); U_s(\theta_B, e = 0) = P;$
$U_b(\mathrm{P}; \mathrm{NR}, if\, Q_H) = 1 - P; U_b(\mathrm{P}; \mathrm{NR}, if\, Q_L) = -P;$
$U_b(\mathrm{P}; \mathrm{GR}, if\, Q_H) = 1 - P - C; U_b(\mathrm{P}; \mathrm{BR}, if\, Q_L) = -P - C.$

In this case, a seller needs to consider the effort in his strategy. The payoff for a good seller and buyer is the same as in the model in Section 4.1, and the payoff for a bad seller is lower by $e(1)$, if he chooses to put forth effort.

34. It is possible to assume this probability to be ψ $(0 < \psi < 1)$, and the main result still holds.

If there is no net reporting cost, $C = 0$, all buyers report. When a bad seller does not make an effort, he will get a bad report, BR. If the game repeats T periods, a bad seller will not make an effort in the last period.

A buyer's willingness to pay is $P_{t+1} = \mu_t + (1-\mu_t)\hat{e}_t$, where $\hat{e}_t$ is the buyer's expectation of the seller's effort. The belief that the seller is a good type in period 2 is:

$$\mu_1 = P(\theta_G \mid \text{GR}) = \frac{P(\text{GR} \mid \theta_G)P(\theta_G)}{P(\text{GR} \mid \theta_G)P(\theta_G) + P(\text{GR} \mid \theta_B)P(\theta_B)}$$
$$= \frac{\mu_0}{\mu_0 + (1-\mu_0)e(1)}.$$

In period t, the updated prior of meeting a good seller is:

$$\mu_{t-1} = \frac{\mu_{t-2}}{\mu_{t-2} + (1-\mu_{t-2})e_{t-1}}.$$

In the last period, T, the buyer's willingness to pay is μ_T.

If $T = 2$, the seller's strategy can be $(e(0), e(0))$ or $(e(1), e(0))$, where the first element represents the action in period $t = 1$, and the second represents the action in period $t = 2$. To examine which strategy is right for the seller, we need to calculate the payoffs.

If the bad seller chooses $(e(0), e(0))$, his total payoff over the two periods is $U_s = \mu_0$. If he chooses $(e(1), e(0))$, his total payoff is

$$U_s = P_1 - e(1) + \delta P_2$$
$$= (\mu_0 + (1-\mu_0)\hat{e}_1 - e(1) + \delta \frac{\mu_0}{\mu_0 + (1-\mu_0)\hat{e}_1}$$
$$= 1 - e(1) + \delta\mu_0.$$

where δ is a discount factor to transform the future payoff to the present value.

If $1 - e(1) + \delta\mu_0 > \mu_0$ (i.e., $e(1) < 1 - (1-\delta)\mu_0$), then the bad seller's best strategy is to make an effort in the first period but not in the second period, $(e(1), e(0))$.

For a T-period game, the payoffs to bad type sellers in each period are the following:

At $t = 1$, $V_1 = P_1 + \delta I(e_1)V_2 - e_1(1)$
If $e_1 = 1$, $I(e_1) = 1$, and $V_1 = 1 + \delta V_2 - e(1)$
If $e_1 = 0$, $I(e_1) = 0$, and $V_1 = P_1 = \mu_0 + (1-\mu_0)\hat{e_1} = \mu_0$
At $t = 2$, $V_2 = P_2 + \delta I(e_2)V_3 - e_2(1)$
If $e_2 = 1$, $I(e_2) = 1$, and $V_2 = 1 + \delta V_3 - e(1)$
If $e_2 = 0$, $I(e_2) = 0$, and $V_2 = P_2 = \mu_1 + (1-\mu_1)\hat{e_2} = \mu_1 = \mu_0$
...
At $t = T-1$, $V_{T-1} = P_{T-1} + \delta I(e_{T-1})V_T - e_{T-1}(1)$

If $e_T = 1$, $I(T-1) = 1$, and $V_{T-1} = 1 + \delta V_T - e(1)$
If $T - 1 = 0$, $I(T-1) = 0$, and $V_{T-1} = P_{T-1} = \mu_0$
At $t = T$, $V_T = P_T = \mu_{T-1} = \mu_0$.

To induce the bad sellers to choose $e_t = 1$ for every period prior to T, the condition $e(1) < 1 - (1-\delta)\mu_0$ must be satisfied. As long as $e(1) < 1 - (1-\delta)\mu_0$, bad sellers will make a genuine effort for $t = 0$ to $t = T - 1$ but will cease to do so in the last period.

Proposition 5.1. *If the reporting cost* $C < 1 - (1-\delta)\mu_0 - e(1)$ *and* $e(1) < 1 - (1-\delta)\mu_0$, *then a bad seller will choose the rebate option and make an effort in the first* $T - 1$ *periods but not in the last period* T.

Proof. See Appendix.

The intuition behind this proposition is similar to that of Proposition 4.1. If the reporting cost is too high, then it is not profitable for a bad seller to put forth an effort and choose the rebate. If the reporting cost is low, then a bad seller will choose the rebate option and also make an effort to provide a high-quality transaction.

Thus, the incentive mechanism can help induce bad sellers to cooperate, helping to sustain a trustful trading environment.

5.2 Types of Rebate Mechanisms and Implementation

Dellarocas et al. (2006) provide an in-depth review of topics related to reputation mechanism design. They point out that the bilateral feedback mechanism adopted by eBay may cause an adverse effect because it allows for retaliation and reciprocation. Even after an unsatisfactory transaction, a buyer who values her own reputation might be reluctant to post "negative" feedback (first) for the fear that the seller might retaliate and leave a negative feedback in response. Reciprocation compromises the effectiveness of feedback mechanisms by unfairly inflating the reputation of users. Klein et al. (2005) provide strong empirical evidence to support this claim. This retaliation problem has drawn attention from academia as well as from eBay. In May 2008, eBay implemented a policy to ban sellers leaving negative feedback to buyers. However, eBay's mandatory policy may lower the market activities as well as feedback participation rates because some sellers switch to other auction sites, and many sellers stop even leaving any feedback to buyers.

In this paper, I propose a mechanism whereby sellers can provide rebates (not necessarily in monetary form) to buyers contingent upon buyers' provision of reports. As mentioned in Section 4.3, this option mechanism is better than requiring all sellers to provide rebates because it allows both bad sellers and good sellers to coexist in the market, but buyers can learn sellers' reputations and bid accordingly.

If the only reason for buyers not to leave negative feedback is the fear of retaliation as Li (2008) suggested,[35] one possible type of rebate mechanism would have sellers commit to setting up an automatic feedback option contingent on receiving payment. In this system, the market maker will leave feedback on behalf of sellers contingent on receiving payments, such that the sellers cannot retaliate against the buyers afterward. In fact, eBay's mandatory policy in May 2008 is a special case of the rebate in the form of automatic feedback. If the fear of retaliation is the only cost of reporting and the cost of fear of retaliation c is less than $(1 - \mu_0)(\alpha - \beta)$ in the whole market, then good sellers will choose the automatic feedback and give up their rights to retaliate against buyers, and bad sellers will mimic good sellers until their types are identified through reports. In this case, eBay's mandatory policy has the same effect as the automatic feedback option when bad sellers mimic good sellers and all buyers pay on time. However, in spite of the benefit of reducing the fear of retaliation in the market, the mandatory policy may have two side effects on feedback system: some sellers will leave eBay, and the participation rate on the feedback system may be lower than before, causing sellers to leave feedback less frequently. When bad sellers expect loss from mimicking good sellers in the long run, they will experience negative payoff under eBay's mandatory policy. As a consequence, they would rather leave the market. This may explain why some sellers switched to other auction sites after eBay announced the policy change in January. Banning sellers from leaving (negative) feedback may hurt sellers' feelings and reduce their interest in participating the feedback system, or sellers will strategically withhold the positive feedback to avoid "rewarding" the buyer who leaves negative feedback. For example, after eBay's announcement many sellers stopped leaving any feedback, and some sellers stopped leaving prompt feedback and waited for buyers to leave feedback first.[36] If the market maker uses the mechanism suggested in this paper, first, we will see feedback for buyers from the automatic feedback; second, buyers will leave feedback without worrying about being retaliated

35. Using data collected from eBay's website, Li (2008) finds that the fear of retaliation may be an important motivation for buyers not to leave (negative) feedback, although the time and effort cost of reporting may not be. One limitation of Li (2008) is that the researcher cannot observe true transaction outcomes from the data set collected on eBay's web site, so the result is subject to the assumption of the independence of error terms in the empirical model.

36. See *http://www.newsday.com/business/ny-bzmain0222,0,6735011.story* (accessed on September 28, 2008.)

against. Therefore, we will see feedback for both buyers and sellers if sellers choose the option.

If the cost of reporting is not only due to the fear of retaliation, and c is not less than $(1 - \mu_0)(\alpha - \beta)$ for some sellers, then, it might not be the social optimal to require every seller to choose the option. This paper assumes that buyers report honestly because most transactions on eBay are one-time events. If relaxing this assumption will allow untruthful feedback from buyers, then we might need another form of rebate.

Another type of rebate mechanism is a combination of the automatic feedback option with monetary incentives. For example, a seller can set up an option to automatically leave feedback contingent on receiving payment and automatically provide monetary incentives contingent on receiving feedback, regardless of whether it is positive or negative. The automatic feedback solves the buyers' concerns of retaliation, and the monetary incentives could cover the rest of the reporting costs. To induce truthful reports, the monetary incentives could be calculated by using truth-eliciting mechanisms.[37] If the rewards in those mechanisms can exceed the total net reporting cost C minus the cost of fear of retaliation, and the reporting cost is less than the threshold in proposition 4.2, sellers will choose the combination of automatic feedback and monetary incentives option as suggested in this paper. Because the amount of monetary incentives depends on how truthful his reports are, the buyer who is willing to tell the truth will get higher monetary incentives as return, thus he will bid higher than the buyer who is going to lie; therefore, the truth-telling buyer will win the auction and report truthfully. As shown in the proof of proposition 4.2 in the appendix, as long as the rebate is higher than the cost, the buyer will incorporate the rebate into his bid. Therefore, only the reporting cost is transferred from the buyer to the seller. The exact way to combine this paper's mechanism with the truth-eliciting mechanism still needs to be cautiously examined by market makers or researchers. One potential problem is that monetary incentives may not motivate people to report or may have a different effect on positive and negative reports, and buyers' feedback may be affected by both the transaction outcome as well as the fact that sellers are paying for it.

With advanced information technology, this mechanism can be easily implemented in a new auction market. The market maker (e.g., eBay) can provide the rebate mechanism options at very low operation costs. However, if eBay or other established auction markets want to implement the mechanism in the existing feedback forum, the "real cost" may be much more. As discussed in Brown and Morgan (2006),

37. See Miller et al. (2005), Jurca and Faltings (2005, 2006, 2007), and Papaioannou and Stamoulis (2005).

a feedback system reform could damage the loyalty of existing users, who might feel that their businesses were harmed by such a change because they relied substantially on eBay's existing rules in determining their business strategies. One possible way to implement the rebate system in the existing feedback system is as follows: eBay could provide an additional feedback table on a user's feedback profile page that shows the feedback after implementing the new rebate system. Then, buyers could observe the seller's past history before implementing the rebate system as well as the seller's reputation after it. When the additional feedback table is first introduced, all the feedback shown in the additional feedback table will play the same role as the inputs in the "past month" feedback in the original feedback table in traders' profile but without missing feedback. If a buyer sees that a seller chooses the rebate option and has the additional table in his profile, then the buyer can compute the posterior probability of a high-quality transaction with the seller and bid accordingly. This is one suggestion on implementation; the more detailed implementation plan maybe be carried by the market makers or determined by future research."

6. Conclusion

The proliferation of online markets means that successful new ideas sometimes have the potential to create a giant. eBay is the largest online auction site in the United State by far, but if eBay's competitors—Amazon, Yahoo!, and Taobao—improve their reputation systems, it is possible that eBay could lose market share.[38] A trustworthy trading environment is essential to the success of online markets. In current online markets, buyers lack incentive to report on sellers' quality. In addition, there is also the problem of buyers' reluctance to report negative feedback for fear of sellers' retaliation. These two conditions have created an environment where lack of adequate information allows bad sellers to commit fraud. To gain and sustain consumers' confidence in the online market, eBay implemented a policy to ban sellers leaving negative feedback to buyers in May 2008. However, eBay's mandatory policy may lower the market activities by driving out some sellers or may discourage sellers to leave feedback.

The goal of this paper has been to introduce ways in which sellers could signal quality and lower transaction costs created by information

38. A study released earlier in May 2006 by the China Internet Network Information Center (CNNIC) confirms Taobao's claims to market dominance in China. Taobao's market share is at 67.3% compared with eBay's 29.1%. From August 2005, Alibaba, the company that owns Taobao, controlled Yahoo's operations in China.

asymmetry.[39] This paper enhances the existing literature on truthful feedback in several important ways. First, I have shown that giving sellers the *option* to compensate buyers for reporting solves the problem of who should and would pay these information costs in a manner that is equitable and efficient. In short, it allows the market to coordinate these costs. Second, I have introduced a mechanism to correct the information asymmetry problem caused by positive feedback bias in online auctions, so that good sellers will be more apt to remain. Unlike eBay's mandatory policy, the mechanism suggested in this paper allows both good sellers and bad sellers to stay in the market, but buyers can distinguish between them through feedback. Third, because the proposed solution is neither coercive nor punitive, it does not require additional regulation or monitoring by market makers, thus minimizing their costs. Fourth, this paper also provides suggestions on the forms of rebates that market makers might choose to use.

In addition to online auctions, this rebate mechanism could also be used for other online markets, such as online retail (e.g., Resellersrating.com and Pricegrabber.com) or price comparison sites (e.g., Shoppers.com, Kelkoo.com, and Nextag.com.) Because there are many sellers in these markets and ratings affect the probability of sales, sellers could also use a rebate option to induce more customer feedback. Another potential type of rebate mechanism is a combination of automatic feedback and monetary incentives. For those in an online retail market, monetary incentives or community recognition might be preferable.

Because the paper makes contributions to the literature from a theoretical standpoint, practitioners should be cautioned about implementing such prescriptions without further development and testing. Potential extensions of this paper include identifying the true reporting cost for buyers to determine the right forms of rebates, examining the strategic timing decisions of traders in bilateral feedback systems, and conducting experiments to test whether buyers' and sellers' behaviors will change as theory predicts.

Appendix A: Proofs

A.1 Proposition 4.1

Proof. First, let us examine the separating equilibrium where good sellers choose rebates (R_S), and bad sellers choose no rebates (NR_S).

39. There are already examples of sellers' signaling behaviors in online auctions. For example, sellers of used Louis Vuitton handbags can acquire certificates that require them to demonstrate a certain level of quality to a third-party evaluator at http://mypoupette.com/.

If it is an equilibrium, then buyers can identify the seller's type by observing whether the seller chooses the rebate option. If the seller chooses it, then she is a good seller, the buyer's willingness to pay is $\alpha + r - C$, and the good seller's payoff $\alpha - C$. Because buyers bid for the product, and rebate on the report is more than the cost of the report, $r > C$, the winning bidder will take the rebate and reporting cost into account when he bids. Otherwise the buyer cannot win the bid. Because the winning bidder includes the rebate in the bidding price, he would choose to report after the transaction. Otherwise he would lose the amount of rebate for which he has paid in the bidding price. If a seller does not choose the rebate option, then she is a bad seller, the buyer's willingness to pay is β, and bad seller's payoff is β. If the reporting cost is larger than the price difference between good and bad sellers, that is, $C > \alpha - \beta$, then both good and bad sellers choose not to rebate (NR_S). If the reporting cost is less than the price difference between good and bad sellers, that is, $C \leq \alpha - \beta$, we need to check whether any sellers want to deviate from the separating equilibrium. A bad seller would get the higher payoff, $\alpha - C$ instead of β, if she pretends to be a good seller by choosing the rebate option. Thus, the separating equilibrium does not exist.

Another separating equilibrium, where good sellers choose not to relate and bad sellers choose rebate, does not exist either. The payoff to the good seller is α, the payoff to the bad seller is $\beta - C$, and $\alpha > \beta$, so that the bad seller can have a higher payoff if he presents as a good seller by choosing a rebate. The bad sellers have incentives to deviate from this separating equilibrium. Another way to check the existence of separating equilibrium is by checking the single-crossing property. Because there is no single-crossing property, that is, the rebate costs the same for the both type sellers, there exists no separating equilibrium. □

A.2 Proposition 4.2

Proof. First, I examine the pooling equilibrium where both types of sellers choose to provide rebates, R_S. In this case, buyers cannot update their beliefs by observing the sellers' choice of providing a rebate. Because both types of sellers provide rebates, all buyers will provide reports. The future buyer can, by using the information about a seller's previous history, update her beliefs on the seller's type. If the buyer does not report, her willingness to pay at period $t + 1$ is

$$P_{t+1} = \mu_t \alpha + (1 - \mu_t)\beta, \tag{A1}$$

although if she chooses to report, her willingness to pay is

$$P_{t+1} = \mu_t \alpha + (1 - \mu_t)\beta - C + r. \tag{A2}$$

The bidding price in period $t+1$ is

$$P_{t+1} = \mu_t\alpha + (1-\mu_t)\beta + r - C. \tag{A3}$$

The payoff for the seller at period t is

$$\begin{aligned} &U_\mathrm{s}(\text{Accept},\ R_\mathrm{S},\ \text{if buyer chooses GR or BR}) \\ &\quad = P - r = \mu_t\alpha + (1-\mu_t)\beta - C. \end{aligned} \tag{A4}$$

It is less than in the case of the benchmark model without the reporting cost in equation (3), and the reporting cost is transferred to the sellers.

If the transaction is high quality, the buyer's payoffs in period t is

$$\begin{aligned} &U_\mathrm{b}(\text{P},\ \text{GR},\ \text{if } Q_\mathrm{H} \text{ and seller chooses } R_\mathrm{S}) \\ &\quad = 1 - P - C + r \\ &\quad = 1 - (\mu_t\alpha + (1-\mu_t)\beta + r - C) - C + r \\ &\quad = 1 - [\mu_t\alpha + (1-\mu_t)\beta]. \end{aligned} \tag{A5}$$

If the transaction is low quality, then it is

$$\begin{aligned} &U_\mathrm{b}(\text{P},\ \text{BR},\ \text{if } Q_\mathrm{L} \text{ andsellerchooses } R_\mathrm{S}) \\ &\quad = -P - C + r \\ &\quad = -(\mu_t\alpha + (1-\mu_t)\beta + r - C) - C + r \\ &\quad = -[\mu_t\alpha + (1-\mu_t)\beta]. \end{aligned} \tag{A6}$$

They are the same as in the case of benchmark model without the reporting cost as in equations (4) and (5). Thus, sellers bear all the reporting cost if the incentive mechanism is adopted.

In the model of this section, I do not allow sellers to change their IDs. As the number of time periods t becomes infinite, according to Weak Law of Large Number, μ for the good seller equals to 1 with very high probability, and the buyer's willingness to pay converges to α; whereas μ for the bad seller equals to to 0 with very high probability, and the buyer's willingness to pay converges to β. If the payoff from providing rebates is higher than the payoff from not providing the rebate in the long run, that is, $\alpha - C \geq P_1 = \mu_0\alpha + (1-\mu_0)\beta$, then the good seller will choose to provide the rebate if he is patient. If the payoff from providing the rebate is higher than the payoff from not providing at the beginning, that is, $P_1 = \mu_0\alpha + (1-\mu_0)\beta > \beta + C$, then the bad seller wants to mimic good seller and choose the rebate until his payoff is less than one identified with a bad seller, that is, $P_t = \mu_t\alpha + (1-\mu_t)\beta < \beta + C$. So as long as $C \leq (1-\mu_0)(\alpha-\beta)$ and $C < \mu_0(\alpha-\beta)$, the patient good sellers will choose to give the rebate, and the patient bad sellers will also give the rebate until their payoff $\mu_t\alpha + (1-\mu_t)\beta - C$ is less than β and choose no rebate otherwise. If $C > (1-\mu_0)(\alpha-\beta)$, both types of sellers want to choose no rebate, $\mathrm{NR_S}$.

Another pooling equilibrium is that both types of sellers choose not to rebate, ($\mathrm{NR_S}$), supported by the off-equilibrium path belief that anyone who chooses to rebate, R_S, must be a bad seller. In this case, the buyer's willingness to pay is the same for all the periods, $P_i = P_1 = \mu_0\alpha + (1-\mu_0)\beta$. Seller's payoff is $P_1 = \mu_0\alpha + (1-\mu_0)\beta$ for every period.

If $C \leq (1-\mu_0)(\alpha-\beta)$, this equilibrium does not exist if we use the intuition criteria. Because the good sellers want to separate from the bad sellers, good sellers have an incentive to give rebates, thus making the buyers report. So the off-equilibrium belief, where anyone who chooses rebate is bad, is not feasible. If $\alpha - C < \mu_0\alpha + (1-\mu_0)\beta$, that is, $C > (1-\mu_0)(\alpha-\beta)$, then the pooling equilibrium in which both types of sellers choose no rebate, $\mathrm{NR_S}$, exists. □

A.3 Proposition 5.1

Proof. If there are reporting costs, then no buyer will be inclined to report. In this case, the buyer's willingness to pay is $P_i = \mu_0$. The good sellers will be worse off than in the case where there is no reporting cost, and the bad sellers will not make an effort in any period. If we use the incentive mechanism proposed in the pure adverse selection model, both types of sellers will choose to give a rebate if the reporting cost is less than the price difference for good transactions and bad transactions, that is, $C < P_{T-1} + \delta I(e_{T-1})V_T - e_{T-1}(1) - \mu_0$. In equilibrium, $P_{T-1} = 1$, $I(e_{T-1}) = 1$, $V_T = \mu_0$, and $e_{T-1}(1) = 1$, the earlier condition is simplified to $1 - (1-\delta)\mu_0 - e(1)$. Also, bad sellers would put forth effort as long as their payoffs are more than μ_0 for each period. □

References

Akerlof, G.A., 1970, "The Market for Lemons: Quality, Uncertainty and Market Mechanism," *Quarterly Journal of Economics*, 84(3), 488–500.

Anderson, K.B., 2005, "Internet Auction Fraud: What Can We Learn from Consumer Sentinel Data?" Technical report, Federal Trade Commission.

Ba, S. and P. Pavlou, 2002, "Evidence of the Effect of Trust Building Technology in Electronic Markets: Price Premiums and Buyer Behavior," *MIS Quarterly*, 26(3), 243–268.

——, B.A. Whinston, and H. Zhang, 2002, "Building Trust in Online Auction Markets Through an Economic Incentive Mechanism," *Decision Support Systems*, 907.

Bajari, P.L. and A. Hortaçsu, 2004, "Economic Insights from Internet Auctions," *Journal of Economic Literature*, XLII, 457–486.

Battigalli, P. and J. Watson, 1997, "Reputation Refinements with Heterogeneous Beliefs," *Econometrica*, 65(2), 369–374.

Bhattacharjee, R. and A. Goel, 2005, "Avoiding Ballot Stuffing in eBay-like Reputation Systems," *Third Workshop on the Economics of Peer-to-Peer Systems*, Philadelphia, PA.

Bolton, G.E., E. Katok, and A. Ockenfels, 2003, "Bridging the Trust Gap in Electronic Markets: A Strategic Framework for Empirical Study," in E. Akcali, J. Geunes, P. Pardalos, H. Romeijn, and Z.S., eds., *Applicatons of Supply Chain Management and E-Commerce Research in Industry*. Dordrecht, The Netherlands: Kluwer Academic Publishers.

Brown, J. and J. Morgan, 2006, "Reputation in Online Auctions: The Market for Trust," *California Management Review*, 49(1), 61–81.

Cabral, L. 2005, "The Economics of Reputation and Trust," Preliminary draft, http://pages.stern.nyu.edu/~lcabral/reputation/reputation.html.

—— and A. Hortaçsu, 2006, "The Dynamics of Seller Reputation: Theory and Evidence from eBay," Working paper http://pages.stern.nyu.edu/lcabral/papers/ebay.pdf.

Celentani, M. and W. Pesendorfer, 1996, "Reputation in Dynamic Games," *Journal of Economic Theory*, 70(1), 109–132.

Cho, I. and D. Kreps, 1987, "Signaling Games and Stable Equilibria," *Quarterly Journal of Economics*, 102(179–221).

Dellarocas, C., 2003a, "The Digitization of Word-of-Mouth: Promise and Challenges of Online Feedback Mechanisms," *Management Science*, 49(10), 1407–1424.

——, 2003b, "Efficiency Through Feedback-Contingent Fees and Rewards in Auction Marketplaces with Adverse Selection and Moral Hazard." In *Proceedings of the 4th ACM Conference on Electronic Commerce*, San Diego, CA, USA.

——, 2004, "Building Trust On-line: The Design of Robust Reputation Mechanisms for Online Trading Communities," in M.N. Doukidis and G.N. Pouloudi, eds., *Information Society or Information Economy? A Combined Perspective on the Digital Era*, Number 95–113. Hershey, PA: Idea Book.

——, 2005, "Reputation Mechanism Design in Online Trading Environments with Pure Moral Hazard," *Information Systems Research*, 16(2), 209–230.

—— and C.A. Wood, 2008, "The Sound of Silence in Online Feedback: Estimating Trading Risks in the Presence of Reporting Bias," *Management Science*, 54, 460–476.

——, F. Dini, and G. Spagnolo, 2006, "Designing Reputation (Feedback) Mechanisms," in N. Dimitri, G. Piga, and G. Spagnolo, eds. *Handbook of Procurement*. Cambridge: Cambridge University Press.

Diamond, D., 1989, "Reputation Acquisition in Debt Markets," *Journal of Political Economy*, 97, 828–862.

Friedman, E. and P. Resnick, 2001, "The Social Cost of Cheap Pseudonyms," *Journal of Economics and Management Strategy*, 10(2), 173–199.

Fudenberg, D. and D.K. Levine, 1992, "Maintaining a Reputation when Strategies are Imperfectly Observed," *Review of Economic Studies*, 59(3), 561–579.

Gazzale, R., 2005, "Giving Gossips Their Due: Information Provision in Games with Private Monitoring," http://econpapers.repec.org/paper/wpawuwpga/0508002.htm.

Houser, D. and J. Wooders, 2006, "Reputation in Auctions: Theory and Evidence from eBay," *Journal of Economics and Management Strategy*, 15(2), 353–369.

Jøsang, A., R. Ismail, and C. Boyd, 2007, "A Survey of Trust and Reputation Systems for Online Service Provision," *Decision Support Systems*, 43(2), 618–644.

Jurca, R. and B. Faltings, 2005, "Enforcing Trustful Strategies in Incentive Compatible Reputation Mechansims," in *In Internet and Network Economics*, volume 3828 of LNCS.

—— and ——, 2006, "Minimum Payments that Reward Honest Reputation Feedback," in *Proceedings of the ACM Conference on E-Commerce*, Ann Arbor, Michigan.

—— and ——, 2007, "Collusion Resistant, Incentive Compatible Feedback Payments," in *Proceedings of the ACM Conference on E-Commerce (EC'07)*, pp. 200–209, San Diego, CA.

Kandori, M., 1992, "Social Norms and Community Enforcement," *Review of Economics Studies*, 69, 63–80.

Kauffman, R.J. and C.A. Wood, 2000, "Running up the Bid: Modeling Seller Opportunism in Internet Auctions," in *Proceedings of the 2000 Americas Conference (AMCIS)*, Long Beach, California.

Klein, B. and K.B. Leffler, 1981, "The Role of Market Forces in Assuring Contractual Performance," *Journal of Political Economy*, 89, 615–641.

Klein, T.J., C. Lambertz, G. Spagnolo, and K.O. stahl, 2005, "Last Minute Feedback," SFB/TR 15 Discussion paper (September 29, 2005).

Kreps, D.M. and R. Wilson, 1982, "Reputation and Imperfect Information," *Journal of Economic Theory*, 27(2), 253–279.

Levine, D.K. and C. Martinelli, 1998, "Reputation with Noisy Precommitment," *Journal of Economic Theory*, 78(1), 55–78.

Li, L.I., 2008, "What is the Cost of Venting? Evidence from eBay," Working paper, Department of Economics, College of Business, University of Louisville.

Milgrom, P. and J. Roberts, 1986, "Price and Advertising Signals of Product Quality," *Journal of Political Economy*, 94(4), 796–821.

Miller, N., P. Resnick, and R. Zeckhauser, 2005, "Eliciting Honest Feedback: The Peer Prediction Method," *Management Science*, 51(9), 1359–1373.

Papaioannou, T.G. and G.D. Stamoulis, 2005, "An Incentives' Mechanism Promoting Truthful Feedback in Peer-to-Peer Systems," in *Proceedings of IEEE/SCM CCGRID.* Washington, DC.

Resnick, P. and R. Zeckhauser, 2002, "Trust among Strangers in Internet Transactions: Empirical Analysis of eBay's Reputation System," *Advances in Applied Microeconomics*, 11, 127–157.

——, ——, J. Swanson, and K. Lockwood, 2006, "The Value of Reputation on eBay: A Controlled Experiment," *Experimental Economics*, 9(2), 79–101.

Saari, D.G., 1999, "Explaining all Three-Alternative Voting Outcomes," *Journal of Economic Theory*, 87, 313–355.

——, 2001, *Chaotic Elections A Mathematician Looks at Voting*. Providence, RI: American Mathematical Society.

—— and K.K. Sieberg, 2001, "The Sum of the Parts can Violate the Whole," *American Political Science Review*, 95(2), 415–433.

Shapiro, C., 1982, "Consumer Information, Product Quality, and Seller Reputation," *The Bell Journal of Economics*, 13(1), 20–35.

——, 1983, "Premium for High Quality Products as Returns to Reputations," *Quarterly Journal of Economics*, 98(4), 659–680.

Xiao, E. and D. Houser, 2005, "Emotion Expression in Human Punishment Behavior," *Proceedings of the National Academy of Science*, 102(20), 7398–7401.