Money Talks: Rebate Mechanisms in Reputation System Design

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

Erte Xiao
Department of Social and Decision Science
Carnegie Mellon University

exiao@andrew.cmu.edu

Abstract:

Reputation systems that rely on voluntary feedback from traders are important in creating and sustaining trust in markets. Feedback is a public good, however, and providing it is often costly. We combine theory with a laboratory experiment to study the effect of a seller precommitment mechanism: sellers have an option to commit by providing a rebate to reduce the buyer's feedback reporting cost before making purchasing decisions. Our theory predicts that this mechanism induces noncooperative sellers to cooperate in the listed-price market. Using a buyer-seller trust game with a unilateral feedback scheme, we find that the seller's rebate decision has a significant impact on the buyer's purchasing decision via *signaling* the seller's cooperative type. More importantly, the experimental data suggest that this precommitment mechanism can significantly improve market efficiency, even though it does not cover the full cost of feedback reporting.

I. Introduction

Asymmetric information is a common problem in exchange environments. The online market is a typical example. Because online traders are usually anonymous and geographically dispersed, participants in these markets often have limited information about their trading partners or products. Efficient exchanges in such markets must consistently rely on trust among traders. Reputation systems as information aggregation devices have been designed to establish and sustain trust.[1] A commonly adopted mechanism to build a reputation system is to rely on voluntary feedback from involved parties. Recent research, however, finds that the feedback systems may not be reliable due to missing or biased information.[2] For example, only about 50% of buyers leave feedback after transactions on eBay. Furthermore, the existing feedback is biased toward the positive due to missing negative feedback.

Previous research has identified two main explanations for the missing or biased information in feedback systems.[3] One is the direct cost associated with reporting feedback (such as time and effort of reporting). The other is the design of the feedback system itself. For example, in an eBay-like bilateral system where both a buyer and a seller can leave feedback for each other after a transaction, the buyer may hesitate to leave negative feedback for fear of seller retaliation. Based on these explanations, various solutions have been proposed to promote contribution in feedback systems. For example, Miller et al. (2005) and Jurca and Faltings (2007) propose truth-eliciting incentive schemes to induce buyers to report honestly. These mechanisms often require the market (e.g., eBay) to provide incentives. Li (2010a) proposes a precommitment mechanism in an online auction market that gives sellers an option to commit by providing a rebate to fully cover the buyer's reporting cost. In theory, this mechanism plays a dual role of providing an incentive for buyers to leave feedback and providing a device for sellers to signal quality or effort to cooperate. However, to apply the theory to practice requires information about the reporting costs for each potential buyer, which is difficult to

---

[1] See Aperjis and Johari, 2010; Archak et al., forthcoming; Bajari and Hortacsu, 2004; Bolton and Ockenfels,2009; Brandts and Figueras, 2003; Brown and Morgan, 2006; Dellarocas, 2003a; Dewan and Hsu, 2004; Fehr et al., 2009; Friedman and Resnick, 2001; Houser and Wooders, 2006; Jin and Kato, 2006; Sobel, forthcoming; Shapiro, 1983; Resnick et al., 2006; and Yamagishi and Matsuda, 2002.

[2] See Resnick and Zeckhauser, 2002; Cabral and Hortacsu, 2010; and Dellarocas and Wood, 2008.

[3] See Avery et al., 1999, Bolton et al., 2004, and Chen et al., 2010; Dellarocas and Wood, 2008, Bolton et al., 2009, Masclet and Pénard, 2008; Bolton et al., 2011, and Li, 2010b.

obtain in real-world markets.

In this paper, we first expand the model of the precommitment mechanism to a more common market in the naturally occurring environment—a listed-price market. We also add heterogeneous beliefs and reporting costs when building the model and deriving the hypothesis. We then conduct a lab experiment to examine the effect of the precommitment mechanism on market efficiency by using monetary incentives as the rebate form. In addition to testing our theoretical predictions, the experiment provides valuable information on people's feedback and trading behavior that have important practical implications.

In particular, we design the experiment to answer the following questions: (1) Does the rebate mechanism improve market efficiency by covering the reporting cost or through its signaling effect? (2) Is the honesty of the reports affected by the monetary incentives? (3) Can we eliminate the positive bias of feedback by shutting down the possibility of retaliation as eBay tried to do in May 2008?[4] To help answer the first question, we use a rebate amount that is not sufficient to cover the full report cost in our experiment. To answer the last question, we use a unilateral feedback system where only buyers can leave feedback to sellers. Moreover, sellers and buyers are randomly and anonymously rematched after each round.

The answers to these questions can help guide the design of these mechanisms in the real world. For instance, it is important to differentiate whether the rebate mechanism improves market efficiency by covering the reporting cost or by providing a signaling effect. The first reason suggests the amount of rebate matters, that the more it covers the cost, the more likely buyers will be to leave feedback and the more efficient the reporting system could be. The second reason suggests that the amount of rebate is not the key, rather the seller's willingness to provide a rebate signals that he is a cooperative type in a way that the buyer can identify even if there is no feedback.

We find that in contrast to the positive bias observed in the eBay studies, our feedback is negatively biased because the reporting cost has significant negative effects on the buyer's propensity to leave feedback when the seller cooperates but not when the seller defects. We also find that a seller's rebate offer increases the buyer's propensity to report only when the seller cooperates, not when the seller defects, and the rebate does not affect the honesty of the feedback. Thus, the rebate

---

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

3

mechanism can reduce the bias of feedback. More importantly, we find that sellers' rebate decisions affect buyers' purchasing decisions by signaling sellers' cooperative type rather than by compensating the ex post reporting cost. Consistent with the signal, the more often a seller provides rebates, the more likely she will cooperate. Under the rebate mechanism, market efficiency increases with the frequency of rebates received by the buyers.

This paper contributes to both academic research and mechanism design in the real business world on how to improve reputation systems to solve asymmetric information problems. First, it generalizes the precommitment model of Li (2010a) to a listed-price market with less restrictive assumptions and shows that it can promote voluntary cooperation in markets characterized by asymmetric information.

More importantly, this paper provides empirical evidence that the mechanism improves market efficiency through its signaling effect, rather than by covering the cost. This finding suggests the amount of the rebate is not the key; instead, market efficiency can be improved as long as a rebate option is provided. In addition, the experimental data show that when reporting is costly, the feedback is negatively biased in the unilateral feedback system. This result suggests that changing a bilateral feedback system to a unilateral system might not be sufficient to correct the bias of feedback, rather it might simply lead the bias in another direction. Our experimental data suggest that the rebate mechanism can reduce such negative bias because subjects are more likely to leave feedback when interacting with cooperative sellers and the honesty of the reports is not affected by the monetary incentives. Thus our study provides evidence that a seller precommitment rebate mechanism can solve the asymmetric information problem of the market and improve market efficiency.

II. Related Literature

Given the importance of the reputation mechanism in facilitating trading in markets that have asymmetric information problems, numerous studies have examined the reliability of feedback systems. Studies by Resnick and Zeckhauser (2002), Cabral and Hortacsu (2010), and Dellarocas and Wood (2008) show that the feedback systems suffer from the problems of missing or biased information. Avery et al. (1999), Bolton et al. (2004), and Chen et al. (2010) argue one reason is that feedback is a public good and there is a direct cost such as time and effort associated with reporting

feedback. Thus, the feedback system suffers the free-rider problem prevalent in public goods provision. Another reason that has been widely studied is the indirect cost of reporting caused by the design of the feedback system itself (see Dellarocas and Wood, 2008, Bolton et al., 2009, and Li, 2010b). For example, in an eBay-like bilateral feedback system where both a buyer and a seller can leave feedback for each other after a transaction, the buyer may hesitate to leave negative feedback for fear of seller retaliation (e.g., Masclet and Pénard, 2008, and Bolton et al., 2011).[5]

In view of these identified problems, scholars have proposed various solutions to improve the effectiveness of feedback systems. For example, Ba et al. (2003) and Dellarocas (2003b) suggest additional monitoring systems to generate feedback and induce sellers to cooperate. Miller et al. (2005) and Jurca and Faltings (2007) propose truth-eliciting incentive schemes to induce buyers to report honestly. Still, all these mechanisms either require buyers to bear the reporting cost or require the market (e.g., eBay) to provide incentives. If the reporting cost remains on the buyer's side, buyers might lack the incentive to leave feedback. If the market provides such incentives, it may burden itself because of a huge volume of transactions. Dellarocas and Wood (2008) provide a highly sophisticated computation mechanism designed to remedy distortions introduced by reporting bias. Their mechanism is computationally involved for nonresearchers, which creates a barrier for real-world application.

Li (2010a) proposes a precommitment mechanism in which sellers have the option to provide rebates (not necessarily in a monetary form) for feedback, regardless whether the feedback is positive or negative. This mechanism transfers the cost to the seller who might be willing to bear the reporting cost because it helps her signal that she will cooperate and thus makes her more attractive to buyers. Once provided, the rebate also provides an incentive for buyers to leave feedback. However, the theory only examines the effect of the precommitment mechanism in an auction market where the seller has the incentive to provide a rebate because it can increase the selling price. Thus, it remains unclear how the rebate mechanism works in a more common market such as a listed-price market. In addition, the paper only discusses the condition when the rebate can cover the reporting cost

---

[5] Another possible explanation for the positive bias is that buyers might tend not to leave negative feedback if they are uncertain about whether a seller harmed them intentionally (see Blount, 1995; Fehr and Gächter, 2000b; and Houser and Xiao, 2010). For example, the seller might have shipped the product on time but the package was delayed by the post office in the real life eBay transaction.

completely. Thus, it is unclear whether the mechanism works because the incentives motivate the buyers to report or because the seller's choice of rebate signals her cooperative type.

This paper contributes to the existing literature by extending this theory of the precommitment mechanism to a listed-price market and providing experimental evidence to understand how the underlying mechanism improves market efficiency and how it might be implemented in the market.

III. The Rebate Mechanism

In this section, we derive theoretical predictions of the outcomes under the precommitment mechanism in a listed-price market environment in which the price of the product is predetermined and fixed. In this paper, the precommitment option takes the form of a monetary rebate offer. Consider a market with M sellers and N buyers where sellers list the same good (g) in each period. For each period, a buyer is randomly matched with a seller and decides whether to buy the product. Suppose M and N are large, so a buyer will not meet the same seller again.[6] If the buyer decides to buy, the seller then decides whether to cooperate. For simplicity, in the rest of the paper, we consider the seller as cooperative when he ships the product as promised.

Suppose a buyer values the product at V_b, a seller's cost of the product is V_s, and the market price of the product is set at P where $V_s < P < V_b$. The utilities for buyer and seller are as follows:

U_b(buyer doesn't buy) = U_s (buyer doesn't buy) = 0;

U_b(buyer buys, seller ships) = V_b- P; U_s(buyer buys, seller ships) = P-Vs;

U_b(buyer buys, seller doesn't ship) = -P; U_s(buyer buys, seller doesn't ship) = P.

We consider a mixed model with both adverse selection and moral hazard with finite T periods.[7] In the model, whether a seller ships is determined by nature according to his type, but he can take unobserved actions to affect the quality of transaction outcomes (to ship or not to ship the product in this case). Suppose there exist two types of sellers: a good type θ_G, who has a high probability α of shipping by nature, and a bad type θ_B, who has a lower probability β of doing so, where $0 \leq \beta \leq \alpha \leq 1$. Suppose also that sellers behave according to their types. For simplicity, we assume that $Vs = 0$ for both types of sellers, and $V_b = 1$ for buyers. Assume that both buyers and

---

6 Resnick and Zeckhauser (2002) report that 89% of all buyer-seller pairs conducted just one transaction during the five-month period covered by their data set.

[7] The model of pure adverse selection is available upon request.

sellers are risk neutral, and $\beta < P < \alpha$. In other words, if buyers can identify sellers' types, they will buy from good sellers, but not from bad sellers.

We consider that buyers may have different experiences or beliefs in online shopping; for example, optimistic buyers may believe that most online sellers are good and pessimistic buyers may believe the opposite. We assume that buyer i has the initial belief that the proportion of good sellers in the market is μ_0^i, and the proportion of bad sellers is $1\text{-}\mu_0^i$.[8] When there is no information about the seller's past history, buyer i's expected payoff is

$$EU_{b0}^i = \mu_0^i \left(\alpha(1 - P) + (1 - \alpha)(-P) \right) + \left(1 - \mu_0^i \right)\left(\beta(1 - P) + (1 - \beta)(-P) \right)$$

$$= \mu_0^i \, \alpha + \left(1 - \mu_0^i \right) \beta - P.$$

If buyer i's willingness to pay, $\mu_0^i \, \alpha + \left(1 - \mu_0^i \right) \beta$, is higher than the listed price P, then buyer i will buy even when there is no reputation information. Let $f\left(\mu^i \right)$ be the probability that buyer i's belief μ^i is greater than $\frac{P - \beta}{\alpha - \beta}$ where $f(\cdot)$ is an increasing function with $f(1) = 1$ and $f(0) = 0$. For a seller, the probability of a successful sale when paired with buyer i is $\Pr(sale) = \Pr(buy) = \Pr$ $(\mu^i \geq \frac{P - \beta}{\alpha - \beta}) = f\left(\mu^i \right)$. Since buyers and sellers are randomly matched in the market, a seller's expected payoff is $EU_{s,0} = E\left(f\left(\mu_0^i \right) \cdot P \right) = P \cdot E\left(f\left(\mu_0^i \right) \right)$ when there is no reputation information. To simplify notations, we denote $E\left(f\left(\mu_0^i \right) \right)$ as $f(\mu_0)$ so $EU_{s,0} = f(\mu_0)P$.

In the naturally occurring environment, sellers usually can decide what quality of transaction outcomes to provide. We incorporate this factor in the model by assuming sellers can make an effort to influence the quality of transaction outcomes and the game repeats T periods. If both good and bad sellers put forth effort ($e = 1$), they will ship the products with probability 1. If they do not put forth effort ($e = 0$), they will not ship the products with probability 1. Assume that for good sellers, the cost of making an effort is 0, $C_{\theta_G}(e = 1) = e(0) = 0$, and for bad sellers, the cost of making an effort is $C_{\theta_B}(e = 1) = e(1) > 0$. To simplify, we assume good sellers will always make an effort, since it costs nothing to them. In this case, a bad seller needs to consider whether to make the effort part of his strategy. A buyer's propensity to buy at period t is $f\left(\mu_{t-1}^i + \left(1 - \mu_{t-1}^i \right)\hat{e}_t \right)$, where $\hat{e}_t$ is the buyer's expectation of a bad seller's effort. The belief that the seller is a good type in period 2 after seeing a good report in period 1 is:

---

[8] Here we allow heterogeneous beliefs of buyers which is not modeled in Li (2010a).

$$\mu_1 = Pr(\theta_G|GR) = \frac{Pr(GR|\theta_G)Pr(\theta_G)}{Pr(GR|\theta_G)Pr(\theta_G) + Pr(GR|\theta_B)Pr(\theta_B)} = \frac{\mu_0^i}{\mu_0^i + (1 - \mu_0^i)e(1)}.$$

Up to period t, if all the past reports are good reports, then the updated prior of meeting a good seller is:

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

Once a seller gets a bad report, then buyers know that he must be a bad seller, so the belief that the seller is a good type equals 0, and a buyer's propensity to buy from him is 0 thereafter. In the last period, T, the buyer's propensity to buy is μ_T, since bad sellers will not make an effort in the last period. Let δ be the seller's discount factor to transform the future payoff to the present value. Based on the model setup, we derive the following propositions.

Proposition 1: *In a complete information market, if $e(1) \leq [1 - (1 - \delta)f(\mu_0)]P$, bad sellers will make a genuine effort for $t = 0$ to $t = T$-1 but will cease to do so in the last period.*

Proof: See Appendix A.

In a complete information market where all buyers report, if the cost of effort is less than the benefit of having a good reputation in the future, a bad seller will choose to make an effort and maintain a good reputation for the T-1 periods. He will not make an effort in the last period, since the last period's reputation will not help him gain from future trade.

Proposition 2: *In an incomplete information market where no one reports, if the rebate cost $0 \leq r < [1 - (1 - \delta)f(\mu_0)]P - e(1)$, then both good and bad sellers will always provide rebates in all periods and in this case, since $e(1) \leq [1 - (1 - \delta)f(\mu_0)]P$, bad sellers will make an effort in the first T-1 periods, but not the last period T.*

Proof: See Appendix B.

No one will report when the reporting cost is more than the maximum internal reporting benefit for all buyers.[9] In this case, bad sellers will have no incentive to make an effort. When the rebate mechanism is introduced, we consider the case that the amount of the rebate is enough to cover

---

[9] There may be psychological benefits from reporting, for example, altruism and reciprocation. Feedback can be a way for buyers to reciprocate, especially when the sellers defect. Previous research suggests that people are often willing to incur costs to punish norm violations or reward good deeds (e.g., Andreoni et al., 2003; Ariely et al., 2009; Bénabou and Tirole, 2006; de Quervain et al., 2004; Fehr and Gächter, 2000a; Nikos, 2008; and Xiao and Houser, 2005).

some buyers' reporting costs so that some buyers are willing to report, and we assume that all buyers report honestly if they decide to report. If a bad seller provides the rebate but does not make an effort, he knows that as long as at least one buyer reports negative feedback his type will be revealed, since good sellers always make an effort to provide good transactions and only bad sellers strategically choose whether to make an effort. Therefore, if a bad seller decides to provide a rebate, his effort decisions will be the same as in the complete information case discussed in Proposition 1.

We find that there is a pooling equilibrium where sellers will choose to give rebates if the rebate cost r is less than the expected benefit from providing the rebate which is equal to the expected payoff of having a good reputation minus the expected payoff when there is no reputation information in the market at all, i.e., $[1 - (1 - \delta)f(\mu_0)]P - e(1)$. As $e(1) = 0$ for the good sellers and $e(1) > 0$ for the bad sellers, the condition $r < [1 - (1 - \delta)f(\mu_0)]P - e(1)$ is stricter for the bad sellers than for the good sellers. When good sellers provide a rebate, bad sellers may or may not provide a rebate depending on the magnitude of $e(1)$. In the last period T, a bad seller has no incentive to make an effort, but he may still offer a rebate because he will not actually incur the cost of the rebate if no one buys from him in the last round. [10]

It is worth emphasizing that, as our theory models demonstrate, this precommitment mechanism introduces the idea of an "inducing" mechanism which has the following features: (1) it provides a precommitment *option* for sellers to signal quality rather than *requiring* them to do so; (2) it allows both good sellers and bad sellers to coexist in the market, but makes it possible for buyers to distinguish between them; thus, (3) it can induce bad sellers to cooperate over time.

IV. Experimental Design

We design four treatments to examine how the feedback reporting cost influences trading behavior and, most importantly, whether the rebate option serves as a signaling device and improves market efficiency. The four treatments are: (1) computer automatic feedback treatment (*Auto_fb*); (2) a higher feedback cost treatment with cost = 10 (*C10*); (3) a lower feedback cost treatment with cost = 5 (*C5*); and (4) the higher feedback cost treatment with a rebate = 5 (*C10r5*). Each treatment consists of ten rounds. In each round, each subject is randomly paired with another subject. One round is randomly

---

[10] This is different from the equilibrium for auction market in Li (2010a) where bad sellers will not choose the rebate in the last round.

selected as the payment round. Each round consists of two stages. The first stage is exactly the same in each treatment. Treatments differ in the second stage.

The first stage is a buyer-seller game (modified from Bolton et al., 2004) described in Figure 1.[11] In this stage, subjects are paired anonymously, with one acting as a seller and the other acting as a buyer. If the buyer decides not to buy the product, the game will end and each participant earns 35E$ (experiment dollars). If the buyer decides to buy the product, then the seller decides whether to ship the product. If the seller ships the product, each participant earns 50E$. If the seller does not ship the product, the seller earns 60E$ and the buyer earns 10E$.[12]

In the second stage, the seller receives feedback. The seller is informed of the feedback he receives in each round. Starting from the second round, each buyer sees the paired seller's feedback from all the previous rounds.

Computer automatic feedback treatments (Auto_fb)

In this treatment, the computer automatically records the feedback for each seller. If the buyer did not buy the product in the first stage, then the seller has no decision to make. In this case, the seller receives "N/A (no report)." If the buyer bought the product and the seller shipped the product, the computer automatically leaves a "+ (positive)" feedback for the seller. If the seller did not ship the product, the seller receives a "- (negative)" feedback. All these are common knowledge. Therefore, in this treatment, the reputation mechanism is complete and perfect in that the costless feedback truthfully and fully reveals the seller's behavior in the past (see Bolton and Ockenfels, 2009).

Feedback cost 10 treatment (C10) and feedback cost 5 treatment (C5).

In both the C10 and C5 treatments (see Figure 2 and 3), after the buyer purchases the product and the seller makes the shipping decision, the buyer can leave feedback for the seller at the cost of 10E$ in the C10 treatment or 5E$ in the C5 treatment.

As we mentioned earlier, if there is some psychological benefit to reporting, buyers might be willing to incur some monetary cost to report. Ex ante, we do not know the magnitude of this value of reporting. We speculate that, for some buyers, the psychological benefit of reporting may be higher than the monetary cost of 5E$ and these buyers will report in this treatment. Our goal is to test the

---

[11] Gazzale and Khopkar (2008) also modified Bolton et al. (2004) to test whether allowing buyers to develop reputations for "information sharing" can increase trust and trustworthiness.

[12] The corresponding payoffs in the model setup are $V_b = 40, V_s = 10,$ and $P = 25$.

effectiveness of the rebate mechanism when the feedback system is ineffective due to the cost of reporting. We design the C10 treatment where the reporting cost is 10E$. This is the highest reporting cost we could impose so that the buyer would not have a negative earning from the game (since when the buyer is cheated, her payoff is 10E$). When we derive the hypotheses below, we assume the reporting cost of 10E$ is high enough that the buyer will not report in this treatment.

If a buyer decides to pay the price to rate her seller, the buyer can leave positive, negative, or neutral feedback. If the buyer decides not to rate the seller, the seller will receive "N/A (no report)" as feedback for that round. In addition, if the buyer does not purchase the product, the paired seller also receives feedback of "N/A (no report)" in that round. That is, if the seller has feedback of "N/A (no report)," this means that either the buyer did not purchase the product or the buyer purchased but did not report feedback. We design the feedback mechanism this way to closely mimic real online markets where there is no feedback, both when there is no trade and when the buyer is not willing to leave feedback.

Feedback cost 10 with rebate 5 treatment (C10r5)

The design of the C10r5 treatment is the same as the C10 treatment except that at the beginning of each round, before the buyer decides whether to purchase the product, the seller first decides whether to provide a rebate of 5E$ to cover half of the reporting cost the buyer will incur if she leaves feedback for the seller (see Figure 4). If the seller provides the rebate and the buyer pays 10E$ to report the feedback in the second stage, there is a 5E$ transfer from the seller to the buyer at the end of the round. On the other hand, as long as the buyer does not leave feedback, the 5E$ rebate transfer does not happen even if the seller provides the rebate. Each buyer can see whether her seller provides the rebate when deciding whether to purchase the product.

We design the rebate value equal to 5E$ so that our treatments allow us to examine the signaling role of rebate. In the C10r5 treatment, if the seller provides a rebate of 5E$, the buyer's payoff scenario is equivalent to that of the C5 treatment. Similarly, if the seller does not provide a rebate of 10E$, the buyer's payoff scenario is equivalent to that of the C10 treatment. Thus, by comparing buyers' decisions in the C5 (or C10) treatment and in their payoff equivalent scenario in

the C10r5 treatment, we can draw inferences regarding the nonmonetary value of sellers' rebate choices to buyers.[13]

Note that the feedback each seller receives is either positive, negative, neutral, or N/A. Since sellers in each treatment have two decisions to make—rebate and ship—then in principle, buyers can leave feedback based on their satisfaction regarding not only the seller's shipping decision but also the rebate decision. For example, buyers might leave negative feedback when the seller did not provide the rebate even if the seller shipped the product. As we report in Section V.A., in this treatment we do not observe any incidence of feedback decisions that are inconsistent with the seller's shipping decision.

Procedure

The experiment was conducted at the Pittsburgh Experimental Economics Laboratory using z-tree (see, Fischbacher, 2007). Subjects were randomly and anonymously assigned a role and the role was fixed in all ten rounds. Each subject was randomly and anonymously rematched with another subject in each round. One round was randomly chosen as the payoff round. We ran 13 sessions in total: three sessions for each of the Auto_fb, C5, and C10 treatments and four sessions for the C10r5 treatment. Each session lasted less than one hour. No one could participate in more than one session. Each subject was paid $5 for showing up in addition to the earnings from the games. The exchange rate was 5E$ = $1. Subjects were paid privately. We obtained observations on 214 subjects: 24 pairs in Auto_fb; 23 pairs in C5; 24 pairs in C10; and 36 pairs in C10r5.

V. Hypotheses

In this section, we derive predictions of sellers' and buyers' decisions based on our theory framework discussed in Section II.

Buyers' reporting decisions

Hypothesis 1: Reporting cost has a negative effect on buyers' reporting behavior:

Pr (report): $0 \cong C10 < C5 < 1$.

---

[13] Note that if we design a rebate value equal to 10E$, when the seller provides a rebate, the buyer's net reporting cost will be zero. To provide a cost equivalent treatment for this condition, we will need another treatment where buyers can decide whether to leave feedback at no monetary cost (which differs from the Auto_fb treatment).

We assume that buyers are risk neutral and utility maximizers, and some buyers'
psychological benefit of reporting may be lower than the monetary cost of 5E\$. Thus, some but not all
buyers will report in this treatment (Pr (report): $0 < C5 < 1$). We also assume the reporting cost of
10E\$ is high enough to overwhelm the psychological benefit of reporting and thus we predict
Pr(report): $0 \cong C10$.

Buyers' purchasing decisions

Hypothesis 2: Buyers are less likely to buy when reporting is costly:

$$Pr(buy): C10 < C5 \cong Auto\_fb.$$

As we mentioned earlier, a good seller always makes an effort, so he will always get positive
feedback, if any. A bad seller may choose to make an effort. If he does not make an effort in some
rounds and the buyers in those rounds report, then he will definitely receive some negative feedback.
As a consequence, no buyer will buy from him in the future. The reason is that good sellers always
make an effort to provide good transactions and only bad sellers strategically choose whether to make
an effort. Thus, as long as at least one buyer reports negative feedback, a bad seller who cares about
future payoff will choose to make an effort.

In the C10 treatment, we assume that the reporting cost is so high that no one reports
feedback. Consequently, bad sellers will not make an effort, and only optimistic buyers will buy. In
the C5 treatment, bad sellers know that if they fail to make an effort, some buyers may leave feedback.
Consequently, they will make an effort for higher future payoffs. Thus, the probability of purchasing
in the C5 treatment is about the same as that in Auto_fb case where all reports are revealed, and both
of these have a greater probability than the C10 treatment.[14]

Hypothesis 3: Buyers take "no rebate" as a signal of a bad seller:

$$Pr^{C10}(buy) > Pr^{C10r5}(buy \mid No\ Rebate) \quad \& \quad Pr^{C10r5}(buy \mid Rebate) > Pr^{C10r5}(buy \mid No\ Rebate).$$

In the theoretical model, both types of sellers will choose to give rebates as long as it is profitable for
bad sellers to mimic good ones. Therefore, providing "rebate" does not provide any information about

---

[14] If our assumption that no one reports in C10 is invalid (i.e., some buyers still report or sellers believe that
some buyers will report), in the pure adverse selection model the result in the C10 treatment would be
indifferent from that in the C5 treatment. In the mixed model case, bad sellers would make an effort to ship the
products like in the C5 and Auto_fb treatments. As a result, the probability of buying in the C10 treatment
would be indifferent from the C5 and Auto_fb treatments.

the type of seller, but providing "no rebate" does. The reporting cost in the C10 treatment is the same as in the C10r5 treatment when the seller does not offer a rebate, but the latter case provides additional information about the seller's quality. Thus, ceteris paribus, buyers should be less likely to buy the product in the latter case than in the former.

Comparing the case in the C10r5 treatment where the seller offers a rebate with the case where the seller does not offer a rebate, buyers expect the seller to be more likely to cooperate in the former than in the latter case. Additionally, the reporting cost for the buyer is lower in the former case than in the latter case. Thus, we predict that in the C10r5 treatment, buyers will be more likely to buy when the seller offers a rebate than when the seller does not.

Sellers' rebate and shipping decisions

Hypothesis 4: Sellers are less likely to ship when reporting is costly:

Pr(ship): $C10 < C5 \cong Auto\_fb.$ In the mixed model, the analysis of sellers' shipping decisions is similar to the analysis in Hypothesis 3 because buyers' buying decisions depend on how they think sellers will behave. In the C10 treatment, we assume that no one will report. Thus, bad sellers will not put forth an effort, and only optimistic buyers will buy. In the C5 treatment, at least some buyers will report. Thus, as in the Auto_fb treatment, bad sellers in C5 treatment will make an effort in order to have future payoffs. As a result, the probability of shipping is higher than in the C10 treatment.[15]

Hypothesis 5: Sellers' decisions not to offer a rebate signals an intention to defect:

$Pr^{C10r5}(\text{ship}| \text{Rebate}) > Pr^{C10r5}(\text{ship} | \text{No Rebate}).$

Our theoretical framework discussed earlier suggests that bad sellers will mimic good sellers by choosing the rebate option if it is profitable for them to continue providing rebates. Although the decision to provide rebates does not help us identify good sellers and bad sellers, a decision not to provide a rebate signals the seller's intention to defect. Therefore, we predict that in the C10r5 treatment, sellers who do not provide a rebate will be less likely to ship the product than those who offer a rebate.

---

[15] If our assumption that no one reports in the C10 treatment is invalid, then the probability of shipping in the C10 treatment would be indifferent from the C5 and Auto_fb treatments for the same reason stated in footnote 14.

From our hypotheses, we expect that under the rebate mechanism, market efficiency is increasing with the frequency of rebates offered by the sellers. We define efficient trades as cases where the buyer buys and the seller ships. As discussed in Hypothesis 3, when the rebate mechanism is introduced, a buyer is more likely to buy when the seller offers a rebate than when the seller does not offer a rebate. Meanwhile, Hypothesis 5 suggests that, under the rebate mechanism, the shipping rate is higher when sellers offer a rebate than when sellers do not offer a rebate. Combining these two hypotheses, we predict that the more often a buyer receives rebates, the more efficient trades the buyer can achieve.

VI. Results

In this section, we first report results of buyers' decisions, followed by results of sellers' decisions. We then examine how the rebate mechanism affects market efficiency.

VI. A. Buyers' feedback reporting decisions

Our data suggest that buyers report feedback honestly most of the time in the C5, C10, and C10r5 treatments. Among the 1070 transactions in all treatments, only three buyers left negative or neutral feedback to a seller who shipped the product.[16] In all the other cases where the buyer chose to leave feedback, she left positive feedback for cooperative sellers and negative feedback for noncooperative sellers. Thus, in the following analysis, we focus only on the frequency of reporting and not the honesty of reporting.

We first compare buyers' feedback reporting behavior in the C5 and C10 treatments to test Hypothesis 1.

Result 1. The cost of feedback has a significant negative effect on the buyer's reporting decision. We calculate the feedback reporting rate among those buyers who purchased the product (see Table 1) in each treatment. Supporting our Hypothesis 1, the cost of feedback has a significant and negative effect on the buyer's reporting decision. About 32% of buyers left feedback for the seller when the cost of doing so was 5E$. In contrast, when the reporting cost was 10E$, only about 10% of buyers

---

[16] One buyer in the C5 treatment left negative feedback in one round after the seller shipped the product. The feedback behavior of this buyer in the following rounds, however, is consistent with the shipping decisions of her paired seller. So it might be just a mistake made by this buyer in the early rounds of the session. One buyer in the C10 treatment left neutral feedback in one round after the seller shipped the product. Also, one buyer in the C10r5 treatment left neutral feedback after her seller provided a rebate and shipped the product.

left feedback. Using the data from Treatment C10 and Treatment C5, we ran a random individual effect logit regression analysis of the probability of buyers who bought the product to leave feedback (see Table 2 column (1)). The independent variables include round, dummy for the final round, and two treatment dummies.[17] The coefficients of the two treatment dummies are significantly different (chi-square test, p=0.01).[18]

Previous research has shown that people are more willing to punish bad behavior than to reward good behavior (e.g. Keysar, et al. 2008). If reporting feedback is a way for the buyer to reciprocate to the seller, we should expect that buyers would be more willing to leave negative feedback when the seller did not ship the product than to leave positive feedback when the seller shipped the product. In light of this, we next investigate how buyers left feedback in the C5 and C10 treatments and compare how cost affects a buyer's propensity to leave feedback when the seller ships versus when the seller does not ship the product.

Result 2. *Buyers are more likely to leave feedback when the seller does not ship the product than when the seller ships the product. Moreover, the propensity of buyers to leave feedback is more sensitive to the reporting cost when the seller is cooperative than when the seller defects.*

As shown in Table 1, in both the C5 and C10 treatments, buyers were more likely to leave feedback when the seller did not ship the product than when the seller shipped the product (50% vs. 28 % in the C5 treatment and 40% vs. 4% in the C10 treatment, respectively). Thus, reporting behavior is negatively biased.

To provide statistical evidence, we ran a random individual effect logit regression analysis of feedback reporting behavior (see Table 2column (2)) similar to the regression analysis reported in Table 2 column (1). The only difference is that we separated the cases of cooperative sellers from the noncooperative sellers. Thus, the independent variables include round, last round dummy, interactions of treatment dummies, and whether the seller shipped the product. We find that the coefficient of C10ship is significantly different from that of C10noship and the coefficient of C5ship is significantly different from that of C5noship (p<0.01).

---

[17] In all the regression models reported below, we always first consider the possible session effects by including dummies for each session in the regressions. We find the session dummies are not jointly significant in any regression (chi-square test, p>0.10). Thus, we report regressions without including session dummies.

[18] All the coefficient comparison tests reported henceforth are chi-square tests with two-tail p-values.

Moreover, the coefficient of C10noship is not significantly different from C5noship (p=0.19). However, the coefficient of C10ship is significantly different from C5ship (p<0.01). These results suggest that the propensity of buyers to leave feedback is less sensitive to the reporting cost when the seller defects than when the seller cooperates. Thus, higher reporting costs lead to more negative-biased reporting behavior.[19] This is consistent with previous findings that there exists an asymmetry in positive and negative reciprocity.[20]

Next, we consider the effect of rebates on buyers' reporting behavior.

Result 3. *When the seller cooperates, a buyer is more likely to leave feedback when the seller offers a rebate than when the seller does not offer a rebate.*

Table 1 also reports the feedback reporting rate in each scenario in the C10r5 treatment depending on whether the seller offered a rebate and whether the seller shipped the product. It shows that, when the seller shipped the product, the buyer is more likely to report when the seller offered a rebate than when the seller did not offer a rebate (26% vs. 12%). In contrast, when the seller did not ship the product, the reporting rate is lower when the seller provided a rebate than when the seller did not provide a rebate (39% vs. 62%) although as we will show this difference is not statistically significant.

To provide statistical evidence of the effect of a rebate on buyers' reporting decisions, we expand the random effect logit regression model in Result 2 by including the data from the C10r5 treatment. For C10r5 treatment data, we also separate the cases where the seller cooperated from the cases where the seller defected and the cases where the seller provided a rebate from the cases where the seller did not provide a rebate. The results of the regression are reported in Table 2 column (3). First, we find that the coefficient of C10r5_rebship is significantly different from C10r5_norebship (p<0.01) and C10ship (p<0.01), but not significantly different from C5ship (chi-square test, p= 0.96).

---

[19] One explanation is that, as sellers' decision is binary (ship vs. not ship), buyers who are willing to incur the cost to leave the feedback may implicitly collude and take N/A (no feedback) as an indicator that the seller shipped the product and negative feedback as indicating the seller did not ship. Thus, these buyers can economize on the cost of building a reputation system by just leaving negative feedback when the sellers did not ship and not leaving positive feedback when the sellers shipped. In C10 treatment, the increased cost of reporting may facilitate this collusion and buyers are thus much less likely to leave positive feedback.

[20] See Abeler et al., 2010; Al-Ubaydli and Lee, 2009; Gazzale and Khopkar, 2008; Keysar et al., 2008; Offerman, 2002.

Meanwhile, the coefficient of C10r5_norebship is significantly different from C5ship (p=0.02), but not significantly different from C10ship (p=0.56). These results suggest that rebates, by compensating for the cost, can make buyers more likely to leave feedback when sellers ship the product.

On the other hand, all the coefficients of the independent variables related to the cases where the sellers failed to ship are not significantly different between each other. For example, the coefficient of C10r5_rebnoship is not significantly different from that of C5noship or C10noship or C10r5_norebnoship (p>0.10). The coefficient of C10r5_norebnoship is not significantly different from that of C5noship or C10noship (p>0.10). These results suggest that rebates have no significant effect on buyers' reporting decisions when sellers fail to ship the product.

Thus, our data suggest that *if buyers' reporting behavior is negatively biased then offering a rebate can reduce such negative bias because it increases the reporting rate only when the seller cooperates and the rebate offer does not affect the truth of feedback.*

VI. B. Buyers' purchasing decisions

To test Hypotheses 2 and 3, we first report the comparison of buyers' purchasing behavior in the Auto_fb, C5, and C10 treatments to see how buyers' purchasing behavior is affected by the feedback reporting cost. We then compare buyers' purchasing behavior in the C10r5 treatment with Auto_fb, C5, and C10 to explore how it is affected by the rebate mechanism.

As shown in Table 1, on average, buyers purchase the product 79% of the time in the Auto_fb treatment, 70% in the C5 treatment, and 68% in the C10 treatment. This result is consistent with our Hypothesis 2. However, these differences are not statistically significant. We start our analysis with a simple random individual effect logit regression model of buyers' purchasing decisions, including only round, last round dummy, and three treatment dummies as independent variables. The results are reported in Table 3 column (1). We find that the coefficients of C5, C10, and Auto_fb are not jointly significantly different from each other (p=0.23). Nor is any pairwise comparison of the coefficients of any two treatments significant (p>0.10).

Result 4. *Supporting Hypothesis 3, buyers are significantly less likely to buy the product when the seller did not provide the rebate in the C10r5 treatment than when the seller provided the rebate or when there is no rebate opportunity at all.*

As reported in Table 1, in the C10r5 treatment, buyers purchase the product about 84% of the time if the seller provides a rebate (the highest purchase rate among all the conditions), but only about 36% of the time if the seller does not provide the rebate (the lowest purchase rate among all the conditions). To provide statistical evidence comparing the purchase rates in each condition, we further expand our regression model discussed above (Table 3 column (1) regression) by including data from all four treatments.

First, we add dummies of whether the seller provided a rebate to the buyer in this round ($C10r5\_noreb_{i,t}$ and $C10r5\_reb_{i,t}$). We also expand our regression model by adding the current matched seller's reputation variables and the shipping decision of the matched seller in the previous round. Previous research shows that the reputation history of sellers has a different effect on buyers' decisions depending on when the feedback is reported. In particular, the most recent feedback is more important for the buyer than earlier feedback (see Dellarocas, 2003a). In light of this, we separate the most recent feedback in the previous round from the other feedback received in the past (round 2 to round t-2). For each buyer i in round t, we calculate the total amount of positive feedback and negative feedback received by the matched seller from round 2 to round t-2, respectively.

Thus, the independent variables in the regression include: round (t), final round dummy, C5 treatment dummy (C5), C10 treatment dummy (C10), Automatic feedback treatment dummy (Auto_fb), total number of positive feedback the matched seller received up to round t-2 $\left(\sum_{t=1}^{t-2} Positive\, fb_{i,t}\right)$, total number of negative feedback the seller received up to round t-2 $\left(\sum_{t=1}^{t-2} Negative\, fb_{i,t}\right)$, whether the seller received positive feedback in the previous round $\left(Positive\, fb_{i,t-1}\right)$, whether the seller received negative feedback in the previous round $\left(Negative\, fb_{i,t-1}\right)$, whether the seller shipped the product in the previous round $\left(Ship_{i,t-1}\right)$, and whether the seller provided a rebate to the buyer in this round ($C10r5\_noreb_{i,t}$ and $C10r5\_reb_{i,t}$). The results are reported in Table 3 column (2).

Consistent with previous literature (see the survey paper Dellarocas 2003a), the regression results show that buyers are more likely to buy if the seller received positive feedback in the previous round and less likely to buy if the seller received negative feedback in the previous round (the coefficient of $Positive\, fb_{i,t-1}$ is significantly positive and the coefficient of $Negative\, fb_{i,t-1}$ is significantly negative). Buyers also take into account the total amount of negative feedback received

by the seller in the past when deciding whether to buy. The more negative feedback the seller received two rounds ago, the less likely the buyer will be to buy the product (the coefficient of $\sum_{t=1}^{t-2} Negative\ fb_{i,t}$ is significantly negative). Interestingly, however, the amount of positive feedback received two rounds ago does not have a significant effect on a buyer's purchasing decision (the coefficient of $\sum_{t=1}^{t-2} Positive\ fb_{i,t}$ is positive but not significant).

The regression results support Hypothesis 3: "not providing rebate" provides information about the types of sellers, but "providing rebate" does not. To see this, first note that buyers are more likely to buy when the seller offers the rebate than when the seller does not (the coefficient of *C10r5_noreb*$_{i,t}$ is significantly lower than that of *C10r5_reb*$_{i,t}$, p<0.01). Second, when the seller does not provide the rebate, the buyer is much less likely to buy than when the seller does not have the rebate opportunity at all (the coefficient of *C10r5_noreb*$_{i,t}$ is significantly lower than that of *C10*, p<0.01). On the other hand, providing the rebate does not make buyers more likely to buy than when the rebate mechanism is not available (the coefficient of *C10r5_reb*$_{i,t}$ is not significantly different from that of either *C5* or *C10*, p=0.75).

These findings suggest that the rebate mechanism affects buyers' purchasing decision mainly because of its signaling effect, rather than its role in covering the reporting cost. We next provide evidence that the seller's rebate decision is indeed a credible signal of his cooperativeness.

VI. C. Sellers' rebate and shipping decisions

We first test Hypothesis 4 by investigating how sellers' shipping decisions may differ between treatments due to the reporting cost. Then, to test Hypothesis 5, we investigate how sellers make rebate decisions and how rebate decisions are correlated with shipping decisions.

We report sellers' shipping rates in Table 1. Consistent with Hypothesis 4, the shipping rate is highest among the three non-rebate treatments in the Auto_fb treatment (88%) and lowest in the C10 treatment (81%). The difference, however, is small.

Similar to the analysis of buyers' decisions, to provide statistical evidence, we first conduct a random individual effect logit regression analysis of the sellers' shipping decisions by including only round variables, C10, C5, and Auto_fb treatment dummy variables. This allows us to see whether the sellers' shipping decisions overall are significantly affected by the feedback reporting cost. The regression results are reported in Table 4. We find that the coefficients of the three treatment dummies

are not jointly significantly different from one another (p=0.76). Nor is any pairwise comparison of the coefficients of any two treatments significant (p>0.10). Thus, although our data suggest that the shipping rate is decreasing in the reporting cost, this effect of the reporting cost is not significant.

We next examine how sellers make rebate decisions and how the rebate decision is correlated with sellers' shipping decisions.

Result 5. *Sellers are more likely to offer the rebate if they did not offer the rebate in the previous round and the buyer did not buy the product than if they did not offer the rebate in the previous round and the buyer bought the product.*

In the C10r5 treatment, every seller provided a rebate at least once to the buyer. On average, sellers chose to offer a rebate 75% of the time. About 90% of sellers chose to offer a rebate at least half of the time. The distribution of rebate frequency is plotted in Figure 5.

We have shown above that most buyers in the C10r5 treatment did not purchase if the seller did not offer the rebate. Sellers who did not offer the rebate may learn from their experience that they are less likely to sell their product if they do not provide a rebate and thus learn to offer rebates over time. To examine whether sellers' rebate decisions are indeed affected by buyers' decisions, we calculate sellers' frequency of offering the rebate in each of four cases, depending on whether the seller provided the rebate in the previous round t-1 and whether the buyer bought the product in that round: (1) seller provided the rebate and buyer bought the product in round t-1; (2) seller provided the rebate but buyer did not buy the product in round t-1; (3) seller did not provide the rebate but buyer bought the product in round t-1; and (4) seller did not provide the rebate and buyer did not buy the product in round t-1. We find that if the seller provided a rebate in the previous round, her decision on rebate in the current round is not greatly affected by whether the buyer bought the product or not in the previous round (90% vs. 97%). In contrast, if the seller did not provide a rebate in the previous round, he is much more likely to provide the rebate in the current round if the buyer did not purchase the product than if the buyer purchased the product in the previous round (51% vs. 20%).

To provide statistical evidence on these differences, we ran a random individual logit regression analysis of sellers' rebate decisions in round t. The independent variables include only the four dummies corresponding to each of the four scenarios mentioned above. We find that the coefficients for the first two scenarios when the seller provided the rebate in round t-1 (buyer bought

the product in round t-1 vs. did not buy the product) are not significantly different (p=0.15), but those for the second two scenarios where the seller did not provide the rebate in round t-1 (buyer bought the product in round t-1 vs. did not buy the product) are significantly different (p=0.01).

Furthermore, we see more and more sellers offering the rebate over time. Figure 6 plots the proportion of sellers who offer the rebate from round 1 to round 10. The slope is significantly positive (p=0.01), suggesting that over time sellers learn that providing the rebate can increase their chances of making a profit.

Result 6. *The more often a seller provides rebates, the more likely the seller will ship the product.*

To examine whether sellers' shipping decisions are correlated with their rebate decisions, we calculate each seller's frequency of offering rebates and the frequency of shipping the product when the buyer chooses to buy the product. We find that when sellers provide a rebate no more than half of the time, their shipping rate is about 63%; in contrast, this rate is 82% for sellers who offer a rebate more than half of the time. To test whether the correlation between rebate and shipping decisions is significant, we ran a Tobit regression analysis of seller i's average shipping rates in the C10r5 treatment using the frequency of rebates seller i offered to the matched buyers as the independent variable. We find that the coefficient of this independent variable is significantly positive (p<0.01).

VI. D. Market efficiency

To test the prediction that under the rebate mechanism the market efficiency is increasing with the frequency of rebates offered by the sellers, we first examine the number of efficient trades (i.e., the case where the buyer bought and the seller shipped the product) in each treatment. We then compare the earnings of buyers and sellers for each treatment.

Result 7. *The feedback reporting cost reduces the proportion of efficient trades. Under the rebate mechanism, the number of efficient trades is increasing in the frequency of rebate provided to the buyer.*

We define efficient trades as the cases when the buyer bought and the seller shipped the product. Figure 7 plots the proportion of efficient trades in each treatment. In the C10r5 treatment, we also separate the case where sellers provide the rebate from the case where they do not. Figure 7 shows that the proportion of efficient trades is decreasing in the feedback reporting cost. The number

of efficient trades is highest in the case where the seller offers the rebate in the C10r5 treatment and lowest when the seller fails to offer the rebate in the C10r5 treatment. To provide statistical analysis of efficient trades, we define a variable Eftrade $_{i,t}$ = 1 if the buyer bought and received the product, and Eftrade $_{i,t}$ =0 if the buyer did not buy or bought but the seller failed to ship the product. We then calculate the average number of efficient trades over ten rounds for each buyer:

$Avg\_Eftrade_i = \sum_{t=1}^{10} Eftrade_{i,t} /10.$

To examine the effect of rebates on market efficiency, for each buyer in the C10r5 treatment, we calculated over the ten rounds the average proportion of times she was offered a rebate from a matched seller before deciding whether to purchase the product:

$$Avg\_rebtms_i = \begin{cases} \dfrac{\sum_{t=1}^{10} C10r5\_reb_{i,t}}{10}, & \text{if C10r5 treatment} \\ 0, & \text{if Auto}_{fb}, \text{C5, or C10 treatment} \end{cases}$$

where $C10r5\_reb_{i,t} = 1$ if buyer i was offered a rebate in round t in the C10r5 treatment. We then ran an OLS regression analysis of Avg_Eftrade$_i$. The independent variables include four treatment dummy variables and Avg_rebtms$_i$. The regression results are reported in Table 5. We find that the coefficient of C10 is significantly different from Auto_fb ($p=0.04$), although C5 is not significantly different from Auto_fb ($p=0.12$). This suggests that a high reporting cost can reduce the number of efficient trades in the market. The regression result also reveals that the coefficient of C10r5 is significantly different from those of Auto_fb ($p=0.01$) and C5 ($p=0.05$), but not significantly different from C10 ($p=0.07$). This indicates that under the rebate mechanism, if no rebate is ever offered the number of efficient trades is just the same as when there is no rebate mechanism.

On the other hand, the coefficient of Avg_rebtms$_i$ is significantly positive, which suggests that the more often a buyer receives a rebate from the seller, the more efficient trades the buyer achieves. The magnitude of the coefficient of Avg_rebtms$_i$ also suggests that the rebate mechanism, even though it covers only half of the reporting cost, can potentially generate as many efficient trades as in the Auto_fb treatment (e.g., if a buyer always receives a rebate from the matched seller, that is Avg_rebtms$_i$=1). *This result provides additional evidence that the rebate mechanism works mostly by signaling the seller's quality rather than by covering the reporting cost.*

Result 8. *When the rebate mechanism is available, buyers' and sellers' earnings are increasing in the number of rebates offered by sellers.*

23

Similar to the analysis above, we calculate the average earnings over ten rounds for each buyer and seller. We then run an OLS regression analysis of buyers' earnings using the same independent variables as in the regression in Table 5. The results are shown in Table 6. We find the coefficients of C10 and C5 are significantly lower than that of Auto_fb (p<0.01 in both cases), but not different from one another (p=0.75). This suggests that buyers earn significantly less when reporting feedback is costly than when the reputation can be automatically recorded at no cost to the buyers.

The coefficient of C10r5 is significantly lower than the coefficients of the three other treatments (p<0.01 for all the pair-wise tests). This indicates that if buyers are never offered a rebate, they will earn less than if there is no rebate mechanism (i.e., C10 treatment). On the other hand, the coefficient of Avg_rebtms_i is significantly positive, suggesting that a buyer who receives rebates more often can also earn more. Indeed, the regression results suggest that if a buyer receives rebates every time (i.e., $Avg\_rebtms_i=1$), she can earn almost as much as the buyers in the Auto_fb treatment. If a buyer receives rebates more than 80% of the time (i.e., $Avg\_rebtms_i=0.8$), she can earn more than the buyers in the C5 or C10 treatments. Again, this suggests that the rebate not only compensates a buyer's cost but also enables buyers to make efficient trades.

Using the same regression analysis, we next discuss the effect of the reporting cost and the rebate mechanism on sellers' earnings. The regression results are reported in Table 6. We find that the reporting cost does not significantly affect sellers' earnings. In particular, the coefficient of neither C5 nor C10 is significantly different from Auto_fb (p=0.13 and 0.08, respectively). However, the coefficient of C10r5 is significantly different from C10, C5, and Auto_fb (p=0.056, 0.04, and 0.00, respectively). This suggests that when the rebate mechanism is available, not providing the rebate makes the seller earn even less than if the rebate mechanism is not introduced.

On the other hand, the regression results also suggest that providing a rebate positively affects sellers' earnings (coefficient of Avg_rebtms_i is significantly positive). The positive coefficient of Avg_rebtms_i indicates that if the seller offers rebates 80% of the time, he can earn more than she does in the corresponding no-rebate environment (i.e., the C10 treatment). To earn as much as the sellers in the C5 treatment, the seller needs to provide rebates about 90% of the time. If the seller provides a rebate every time, his earnings will be just slightly lower than in the Auto_fb treatment.

VII. Conclusion

Reputation information is crucial to build cooperation in markets where information is asymmetric. However, reputation systems that rely on voluntary feedback have the problem of missing or biased reports, and thus their effectiveness in promoting trust and trustworthiness in the markets is limited. In this paper, we examine an innovative precommitment mechanism (rebate for feedback) aimed at inducing participation in reputation systems and promoting voluntary cooperation in markets where information is asymmetric. We show, both theoretically and experimentally, that the advantage of this precommitment mechanism is to improve the reputation system and market efficiency without requiring buyers or markets to bear the reporting cost.

In the experiment, we use monetary incentives to examine the effectiveness of the rebate mechanism in a listed-price market with a unilateral feedback system. Our experimental data support theoretical predictions of the effectiveness of rebates. One contribution of our experiment is that it provides information on how the underlying rebate mechanism affects buyers' purchasing behavior and improves market efficiency. We provide evidence that the rebate mechanism can increase market efficiency because of its signaling effect, rather than its role in reimbursing a buyer's cost. We also find that the honesty of buyer reports is not affected by the seller's rebate decision. This provides useful information for implementing the rebate mechanism in the real world, that the amount of rebate is not the key: market efficiency can be improved as long as the seller is willing to provide a rebate. This is good news because in a real market it is often hard to accurately quantify the buyer's reporting cost.

In addition, we find that a buyer's propensity to report is more sensitive to the reporting cost when a seller cooperates than when a seller defects. Thus, the feedback from our experiment is negatively biased (see also Gazzale and Khopkar, 2008). This is in contrast to 99% positive observations from the eBay bilateral feedback system, in which the cost is higher for reporting negative feedback than for positive because of the fear of retaliation. This result also suggests that changing the reporting mechanism from bilateral to unilateral, as eBay tried to implement, does not necessarily reduce the bias and may only lead the bias in another direction. Furthermore, we find a rebate offer increases a buyer's propensity to leave feedback when the seller cooperates, but has no significant effect when the seller defects. Thus, the data suggest that introducing a rebate mechanism

in the unilateral system can reduce the negative bias by increasing the reporting rate when the sellers cooperate.

The precommitment mechanism introduces a way to design feedback systems that allows parties to solve the asymmetric information problem in the market and reach a win-win outcome. The monetary rebate is one example of how to implement such a precommitment mechanism in a listed-price market. It opens the door for future research to analyze different rebate mechanisms and various markets and create field experiments to test the mechanism in the real market.

References

Abeler, Johannes, Juljana Calaki, Kai Andree and Christoph Basek. 2010. "The power of apology". *Economics Letters,* 107, 233-235.

Al-Ubaydli, Omar, and Min Sok Lee. 2009. "An experimental study of asymmetric reciprocity." *Journal of Economic Behavior & Organization*, 72(2): 738-749.

Andreoni, James, William Harbaugh, and Lise Vesterlund. 2003. "The carrot or the stick: Rewards, punishments, and cooperation". *American Economic Review*, 93(3), 893–902.

Aperjis, Christina, and Ramesh Johari. 2010. "Optimal Windows for Aggregating Ratings in Electronic Marketplaces." ***Management Science.*** Vol. 56, No. 5, pp. 864-880.

Archak, Nikolay, Anindya Ghose, and Panagiotis Ipeirotis. Forthcoming. "Deriving the Pricing Power of Product Features by Mining Consumer Reviews," *Management Science***.**

Ariely, Dan, Anat Bracha, and Stephan Meier. 2009. "Doing Good or Doing Well? Image Motivation and Monetary Incentives in Behaving Prosocially." *American Economic Review,* 99(1): 544–55.**Avery, Christopher, Paul Resnick, and Richard Zeckhauser.** 1999. "The Market for Evaluations." *American Economic Review*, 89(3): 564-584.

Ba, Sulin, Andrew B. Whinston, and Han Zhang. 2003. "Building trust in online auction markets through an economic incentive mechanism." *Decision Support Systems*, 35(3): 273-286.

Bajari, Patrick and Ali Hortacsu. 2004. "Economic insights from internet auctions." *Journal of Economic Literature*, XLII:457-486.

Bénabou, Roland, and Jean Tirole. 2006. "Incentives and Prosocial Behavior." *American Economic Review*, 96(5): 1652–1678.

Blount, S. 1995. "When Social Outcomes Aren't Fair: The Effect of Causal Attributions on Preferences." *Organizational Behavior and Human Decision Process* 63 (2), 131-144.

Bolton, Gary E., Ben Greiner, and Axel Ockenfels. 2011. "Engineering Trust – Reciprocity in the Production of Reputation Information." Working Paper Series in Economics, University of Cologne, Department of Economics.

Bolton, Gary E., Elena Katok, and Axel Ockenfels. 2004. "How Effective Are Electronic Reputation Mechanisms? An Experimental Investigation." *Management Science*, 50(11): 1587-1602.

Bolton, Gary E., and Axel Ockenfels. 2009. "The limits of trust," in *eTrust*, K. Cook, C. Snijders, V. Buskens, eds., Russell Sage: New York, P15-36.

Brandts, Jordi, and Neus Figueras. 2003. "An Exploration of Reputation Formation in Experimental Games." *Journal of Economic Behavior and Organization*, 50, 89-115.

Brown, Jennifer, and John Morgan. 2006. "Reputation in Online Markets: The Market for Trust," *California Management Review,* Vol. 49(1), pp. 61-81.

Cabral, Luis M. B., and Ali Hortacsu. 2010. "The dynamics of seller reputation: Theory and evidence from eBay," *Journal of Industrial Economics*, 58(2010): 54-78.

Chen, Yan, Max Harper, Joseph Konstan and Sherry Xin Li. 2010**.** "Social Comparisons and Contributions to Online Communities: A Field Experiment on MovieLens." *American Economic Review***,** September,100（4）， 1358-98.

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

Dellarocas, Chrysanthos. 2003b. "Efficiency through feedback-contigent fees and rewards in auction marketplaces with adverse selection and moral hazard." In *Proceedings of the 4th ACM conference on Electronic commerce*, 11-18. New York: The Association for Computing Machinery.

Dellarocas, Chrysanthos, and Charles A. Wood. 2008. "The Sound of Silence in Online Feedback: Estimating Trading Risks in the Presence of Reporting Bias." *Management Science*, 54(3): 460-476.

Dewan, Sanjeev. and Hsu, Vernon. 2004. "Adverse Selection in Electronic Markets: Evidence from Online Stamp Auctions," *The Journal of Industrial Economics,* Vol. LII, No. 4, pp. 497-516.

de Quervain, Dominique J.-F., Urs Fischbacher,Valerie Treyer, Melanie Schellhammer, Ulrich Schnyder, Alfred Buck, and Ernst Fehr. 2004. "The neural basis of altruistic punishment." *Science* 305, 1254–1258.

Fehr, Ernst, Martin Brown and Christian Zehnder. 2009. "On reputation: A microfoundation of contract enforcement and price rigidity." *The Economic Journal*, 119(March), 333–353.

Fehr, Ernst, and Simon Gächter. 2000a. "Cooperation and Punishment in Public Goods Experiments." American Economic Review, 90(4):980-994.

Fehr, Ernst, and Simon Gächter. 2000b. Fairness and Retaliation: The Economics of Reciprocity. *Journal of Economic Perspective*. 14(3), 159-181.

Fischbacher, Urs. 2007. "z-Tree: Zurich toolbox for ready-made economic experiments." *Experimental Economics*, 10(2): 171-178.

Friedman, Eric, and Paul Resnick. 2001. "The social cost of cheap pseudonyms." *Journal of Economics and Management Strategy,* 10(1) 173-199.

Gazzale, Robert S. and Tapan Khopkar. 2008. "Remain Silent and Ye Shall Suffer: Seller Exploitation of Reticent Buyers in an Experimental Reputation System." Working paper.

Houser, Daniel, and John Wooders. 2006. "Reputation in auctions: Theory and evidence from eBay." *Journal of Economics and Management Strategy*, 15(2):353-369.

Houser, Daniel, and Erte Xiao. 2010. "Inequality seeking punishment." *Economics Letters*, 109(1), 20-23.

Jin, Ginger, and Andrew Kato. 2006. "Price, Quality and Reputation: Evidence from An Online Field Experiment." *RAND Journal of Economics, 37(4)*.

Jurca, Radu, and Boi Faltings. 2007. "Collusion resistant, incentive compatible feedback payments." In *Proceeding of the 8th ACM conference on Electronic commerce*, 200-209. New York: The Association for Computing Machinery.

Keysar, Boaz, Benjamin A. Converse, Jiuwen Wang, and Nicholas Epley. 2008. "Reciprocity Is Not Give and Take: Asymmetric Reciprocity to Positive and Negative Acts." *Psychological Science*, 19(12): 1280-1286.

Li, Lingfang(Ivy). 2010a. "Reputation, Trust, and Rebates: How Online Auction Markets Can Improve Their Feedback Mechanisms." *Journal of Economics and Management Strategy*, 19(2): 303-331.

Li, Lingfang (Ivy). 2010b. "What is the Cost of Venting? Evidence from eBay," *Economics Letters*, 108 (2010) 215–218.

Masclet, David, and Thierry Pénard. 2008. "Is the eBay feedback system really efficient? An experimental study." Working Paper Series, Industrial Economics, Center for Research in Economics and Management, University of Rennes 1, University of Caen and CNRS WP 2008-03.

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

Nikiforakis, Nikos. 2008. "Punishment and counter-punishment in public good games: Can we really govern ourselves?" *Journal of Public Economics*, 92(1): 91-112.

Offerman, Theo. 2002. "Hurting hurts more than helping helps." *European Economic Review*, 46(8): 1423-1437.

Resnick, Paul, and Richard Zeckhauser. 2002. "Trust among strangers in internet transactions: Empirical analysis of eBay's reputation system." *Advances in Applied Microeconomics*, 11: 127-157.

Resnick, P., R. Zeckhauser, J. Swanson, and L. Lockwood. 2006. "The Value of Reputation on eBay: A Controlled Experiment," *Experimental Economics,* 9: 79–101.

Shapiro, Carl. 1983. "Premiums for High Quality Products as Returns to Reputations." *The Quarterly Journal of Economics*, 98(4): 659-680.

Sobel, Joel. Forthcoming. "Signaling Games." In *Encyclopedia of Complexity and System Science*, M. Sotomayor(ed.), Springer.

Xiao, Erte, and Daniel Houser. 2005. "Emotion expression in human punishment behavior." *Proceedings of the National Academy of Sciences of the United States of America*, 102(20): 7398-7401.

Yamagishi, Toshio, and Masafumi Matsuda. 2002. "Improving the Lemons' Market with a Reputation System: An Experimental Study of Internet Auctioning." Working Paper, Hokkaido University.

Figure 1. Seller-Buyer game (Auto_fb treatment)
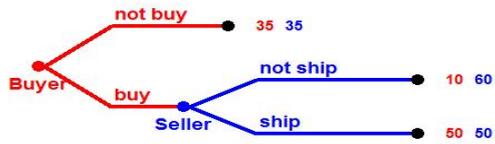
Note: the first number is the buyer's payoff and the second number is the seller's payoff (This applies to fig 2-4 as well).
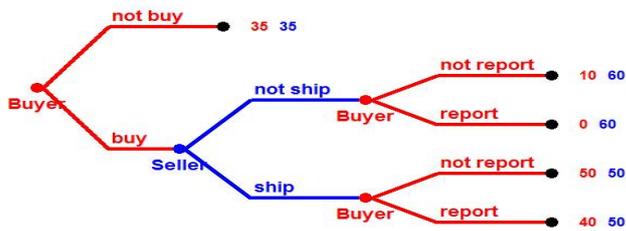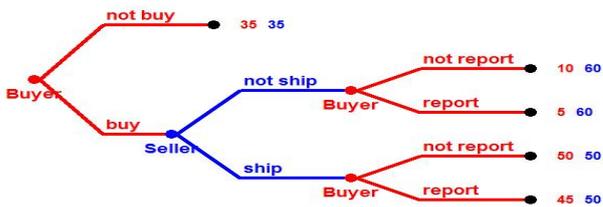
Figure 2. C10 treatment

Figure 3. C5 treatment
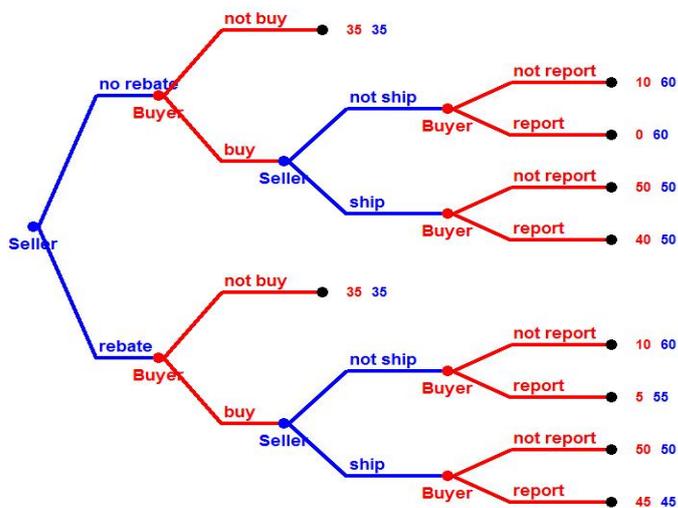

Figure 4. C10r5 treatment

Figure 5. Distribution of sellers' rebate frequency in the C10r5 treatment
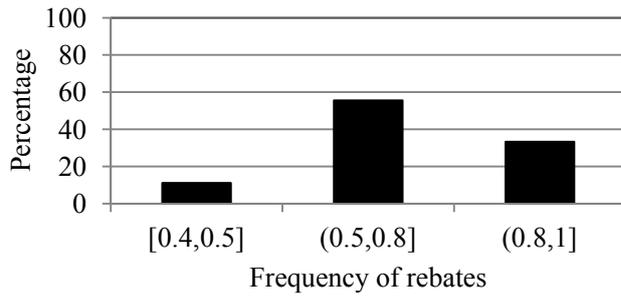

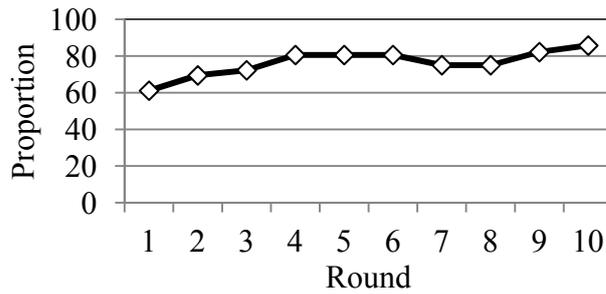**Figure 6.** Proportion of sellers who offered rebates over ten rounds

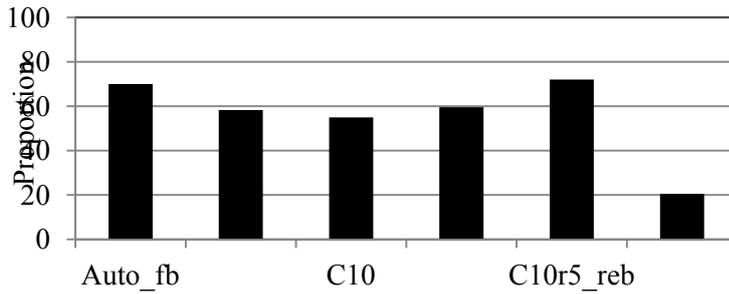**Figure 7.** Proportion of efficient trades in each treatment

Table 1   Descriptive summary of the decisions

Treatment	Purchasing rate (%)	Buyer			Seller
		Reporting rate*			Shipping rate*
		Overall(%)	Seller Shipped (%)	Seller did not ship (%)	
Auto_fb	79.17				88.42
C5	69.57	31.88	28.36	50	83.75
C10	67.5	10.49	3.79	40	81.48
C10r5_reb	83.91	27.85	26.06	38.71	85.84
C10r5_noreb	36.14	33.33	11.76	61.54	56.67

* The rate is calculated based on the cases where buyers chose to buy.

Note: Auto_fb: Computer automatic feedback treatment ; C5: Feedback cost 5 treatment; C10: Feedback cost
10 treatment ; C10r5_reb: the seller provided a rebate in Feedback cost 10 with rebate 5 treatment ;
C10r5_noreb: the seller did not provide a rebate in Feedback cost 10 with rebate 5 treatment

Table 2. Buyers' Feedback Reporting Decisions: Random Individual Effect Logit Regression Model

	Report$_{i,t}$ (=1 if report in round t; =0 if not report)		
	(1) (include C5 and C10 treatments) Coef. (s.e.)	(2) (include C5 and C10 treatments) Coef. (s.e.)	(3) (include C5, C10 and C10r5 treatments) Coef. (s.e.)
Round(t)	-0.09 (0.07)	-0.44 (0.14)	-0.30 (0.08)
Final round dummy	-0.56 (0.90)	-3.69 (1.45)	-2.72 (0.88)
C10	-2.45 (0.67)		
C5	-0.55 (0.60)		
C10ship		-5.17 (1.52)	-4.76 (1.17)
C10noship		2.82 (1.45)	1.84 (1.05)
C5ship		-0.09 (1.06)	-0.49 (0.83)
C5noship		4.85 (1.59)	3.51 (1.09)
C10r5_rebship			-0.54 (0.66)
C10r5_norebship			-3.82 (1.23)
C10r5_rebnoship			2.11 (0.98)
C10r5_norebnoship			1.65 (1.09)
Wald chi2(-)	30.93	27.52	62.02
# of obs.	322	322	571

Table 3. Buyers' Purchasing Decisions: Random Individual Effect Logit Regression Model

	Buy$_{i,t}$ (=1 if purchased in round t; =0 if not)	
	(1) Include Auto_fb, C5 & C10 treatments	(2) Include all the four treatments
Round (t)	-0.04 (0.05)	-0.03 (0.09)
Final round (=1, if t=10)	-1.52 (0.42)	-2.18 (0.56)
C10	1.77 (0.61)	2.77 (0.89)
C5	1.79 (0.60)	2.76 (0.89)
Auto_fb	2.37 (0.59)	1.08 (0.90)
$\sum_{t=1}^{t-2}$ Postive fb$_{i,t}$		0.12(0.13)
$\sum_{t=1}^{t-2}$ Negative fb$_{i,t}$		-1.81(0.35)
Positive $fb_{i,t-1}$		1.56 (0.56)
Negative $fb_{i,t-1}$		-2.75 (0.65)
Ship$_{i,t-1}$		1.31 (0.43)
C10r5_noreb$_{i,t}$		-1.79 (0.87)
C10r5_reb$_{i,t}$		2.51 (0.80)
Wald chi2(-)	45.75	89.56
# of obs.	710	703

Table 4. Sellers' Shipping Decisions: Random Individual Effect Logit Regression
 Model

	Ship$_{i,t}$ (=1 if shipped the product in round t; =0 if not)
	Random effect Logit regression
	Include C10, C5 and Auto_fb treatments Coef. (s.e.)
Round (t)	-0.24 (0.08)
Final round (t=10)	-3.84 (0.83)
C10	4.34 (0.92)
C5	5.02 (1.02)
Auto_fb	4.86 (0.93)
Wald chi2(-)	47.43
# of obs.	512

Table 5. Average proportion of efficient trades: OLS Regression Models

	Avg_Eftrade$_i$
	Include four treatments Coef. (s.e.)
C10	0.55 (0.05)
C5	0.58 (0.05)
Auto_fb	0.70 (0.05)
C10r5	0.13 (0.22)
Avg_rebtms$_i$	0.61(0.28)
R-squared	0.8585
# of obs.	107

Table 6. Average earnings: OLS Regression Models

	Avg_Earning$_i$	
	Buyer	Seller
	Coef. (s.e.)	Coef. (s.e.)
C10	39.42 (0.86)	46.38 (0.57)
C5	39.80 (0.88)	46.57 (0.58)
Auto_fb	43.21 (0.86)	47.79 (0.57)
C10r5	29.26 (3.63)	43.47 (1.39)
Avg_rebtms$_i$	13.67 (4.69)	3.73 (1.72)
R-squared	0.9898	0.9966
# of obs.	107	107

Appendix (Not for publication)

Appendix A: Proof of proposition 1

Proof.

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.

In a complete information market, if the bad seller chooses $(e(0),\ e(0))$, buyer i will buy in period 1 if $\mu_0^i \geq (P - \beta)/(\alpha - \beta)$ but not in period 2 (since the seller does not make effort in period 1). So, the seller's total expected payoff over the two periods is $EU_s = E(f(\mu_0^i) \cdot P) = P \cdot E\left(f(\mu_0^i)\right) = f(\mu_0)P$.

If he chooses $(e(1),\ e(0))$, the buyer will buy in period 1 and 2. So the seller's total expected payoff is:

$EU_s = E[\Pr(sale)_1\, P - e(1) + \delta \Pr(sale)_2\, P] =$

$E[f\left(\mu_0^i + (1 - \mu_0^i)\hat{e}_1\right)P - e(1) + \delta f(\frac{\mu_0^i}{\mu_0^i+(1-\mu_0^i)\hat{e}_1})P]$, where δ is a discount factor to transform the future payoff to the present value. If $\hat{e}_1 = 1$, then $EU_s = E[P - e(1) + \delta f(\mu_0^i)P] = P - e(1) + \delta f(\mu_0)P$. If $P - e(1) + \delta f(\mu_0)P \geq f(\mu_0)P$, i.e., $e(1) \leq [1 - (1 - \delta)f(\mu_0)]P$, 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 expected payoffs to bad type sellers in each period are the following:

At $t=1$, $V_1 = E[\Pr(sale)_1\, P + \delta I(e_1)V_2 - e_1(1)]$

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

$V_1 = E[\Pr(sale)_1 P] = E[f\left(\mu_0^i + (1 - \mu_0^i)\hat{e}_1\right)P] = E(f(\mu_0^i)P) = f(\mu_0)P.$

At $t=2$, $V_2 = E[Pr(sale)_2\, P + \delta I(e_2)V_3 - e_2(1)]$

If $e_2 = 1$, $I(e_2) = 1$, and $V_2 = P + \delta V_3 - e(1)$.

If $e_2 = 0$, $I(e_2) = 0$, and

$V_2 = E[Pr(sale)_2 P] = E[f\left(\mu_1^i + (1 - \mu_1^i)\hat{e}_2\right)P] = E(f(\mu_1^i)P) = E(f(\mu_0^i)P) = f(\mu_0)P.$

………

33

At t=T-1, $V_{T-1} = E[\Pr(sale)_{T-1} P + \delta I(e_{T-1})V_T - e_{T-1}(1)]$

If $e_{T-1} = 1$, $I(e_{T-1}) = 1$, and $V_{T-1} = P + \delta V_T - e(1)$.

If $e_{T-1} = 0$, $I(e_{T-1}) = 0$, and $V_{T-1} = E[\Pr(sale)_{T-1} P] = f(\mu_0)P$.

At t=T, $V_T = E[\Pr(sale)_T P] = f(\mu_0)P$.

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

Appendix B: Proof of Proposition 2

Proof.

If the reporting cost is more than the maximum reporting benefit, then no buyer will be inclined to report. In this case, the seller's expected payoff is $f(\mu_0)P$. 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. When the rebate mechanism is introduced and the amount of rebate is enough to cover some buyers' reporting costs, some buyers are willing to report when provided a rebate. We assume that all buyers report honestly if they decide to report. If a bad seller provides the rebate but does not make an effort, he knows, as long as at least one buyer reports negative feedback, his type will be revealed. The reason is that good sellers always make an effort to provide good transactions and only bad sellers strategically choose whether to make an effort. Therefore, if a bad seller decides to provide a rebate, his effort decisions will be the same as in the complete information case discussed in Proposition 1. Since in this case, $e(1) \leq [1 - (1 - \delta)f(\mu_0)]P$ *is always true,* a bad seller will make a genuine effort from t=1 to t=T–1, but will cease to do so in the last period.

We next show that when we introduce a rebate mechanism, there is a pooling equilibrium where both types of sellers will choose to give rebates if the rebate is less than the expected payoff of having a good reputation minus the expected payoff when there is no reputation information in the market at all,

i.e., $r < E[\Pr(sale)_{T-1} P + \delta I(e_{T-1})V_T - e_{T-1}(1)] - f(\mu_0)P$. In the equilibrium we derived in appendix A, $\Pr(sale)_{T-1} = 1$, $I(e_{T-1}) = 1$, $V_T = f(\mu_0)P$, and $e_{T-1}(1) = e(1)$, the above condition is simplified to $r < [1 - (1 - \delta)f(\mu_0)]P - e(1)$. The reason for the pooling equilibrium is the following: given the off-equilibrium path belief that anyone who chooses not to provide rebates must be a bad seller, if good sellers choose the rebate option but bad sellers do not, then buyers can immediately identify bad sellers by observing which ones do not provide rebates. Therefore, bad sellers have an incentive to mimic good sellers by providing rebates if the expected payoff of doing so is higher than the expected payoff of not doing so — i.e., r < $[1 - (1 - \delta)f(\mu_0)]P - e(1)$. Since we assume that the good sellers' effort cost is 0, good sellers will always provide rebates and make an effort if $r < [1 - (1 - \delta)f(\mu_0)]P$. Thus, (1) if $r < [1 - (1 - \delta)f(\mu_0)]P - e(1)$, both bad and good sellers will provide rebates; (2) if $[1 - (1 - \delta)f(\mu_0)]P - e(1) < r < [1 - (1 - \delta)f(\mu_0)]P$, only good sellers will provide rebates; and (3) if $r > [1 - (1 - \delta)f(\mu_0)]P$, neither good sellers nor bad sellers will give rebates.

As we mentioned above, in the last round, bad sellers will not make an effort. Expecting this, buyers will decide whether to buy according to their initial beliefs. In a listed-price market where the rebate is less than the price *P*, offering a rebate in the last period will not make a bad seller worse than not doing so, since he will not lose the rebate if no one buys from him in the last period. ∎

Appendix C: An Adverse Selection Model

To understand the intuitions behind the rebate mechanism, we use a simple setup of a pure adverse selection model with infinite periods. The basic setup is similar to the mixed model of both adverse selection and moral hazard discussed in Section II, except that sellers' cannot take actions to affect the transaction outcomes in the pure adverse selection model. The transaction outcomes are solely determined by nature according to the seller's type.

If a seller's complete past reputation history is available to buyers, then a buyer can update her belief by using the information. Buyer *i*'s belief of meeting a good seller in period *t* is

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

where t_{GR} and t_{BR} represent the number of good and bad reports the seller received in the past, respectively. A seller's expected payoff at period t is $EU_{st} = Pr\left(\mu_t^i \geq \frac{P-\beta}{\alpha-\beta}\right) \cdot P = f(\mu_t) \cdot P$.

In the first round, a buyer whose initial belief μ_0^i is less than $\frac{P-\beta}{\alpha-\beta}$ does not buy. In the second round, after the buyer observes sellers' reputations, the buyer updates her belief and decides whether to buy based on the value of $\mu_t^i \alpha + (1 - \mu_t^i)\beta$ and P. If a seller earns a good (or bad) reputation in period 1, then in period 2 there will be more (or fewer) buyers whose updated belief is that μ_t^i is greater than $\frac{P-\beta}{\alpha-\beta}$. Since good sellers have a higher probability of shipping than bad sellers, in the long run, as t goes to infinity, $\mu_t = 1$ with high probability for a good seller and $\mu_t = 0$ with high probability for a bad seller.[21] Therefore, in the long run, when a seller's complete reputation is revealed, a good seller has a higher probability of selling than a bad seller.

When there is no information available, only optimistic buyers whose prior belief was $\mu_0^i \geq \frac{P-\beta}{\alpha-\beta}$ will buy. If the true proportion of good sellers in the market is less than $\frac{P-\beta}{\alpha-\beta}$, then optimist buyers will be better off not buying. If the true proportion of good sellers in the market is greater than $\frac{P-\beta}{\alpha-\beta}$, then pessimistic buyers whose prior belief is $\mu_0^i < \frac{P-\beta}{\alpha-\beta}$ will be better off buying. When there is complete information, buyers' purchasing decisions will depend on their updated beliefs about μ_{t-1}. If μ_{t-1} goes up (or down) so that the value of $\mu_{t-1}\alpha + (1 - \mu_{t-1})\beta$ is more (or less) than P, then the buyer will buy (or not buy). Compared with the no-information case, the total buyers' welfare increases when there is complete information. When changing from the no-information case to the complete information case, there is a wealth transfer from bad sellers to good sellers because buyers can identify sellers' types in the complete information case.

If there is a positive reporting cost for all buyers, say C, then future buyers may not update their beliefs correctly. To see this, suppose that buyers' internal reporting benefit is $b \in [0, B]$. If the reporting cost C is higher than the maximum reporting benefit B for all buyers, then no buyer reports.

---

[21] Proof is provided in the Appendix D.

If *C< B*, we consider two cases. In the first case, if the benefits from giving good reports and bad reports are not symmetric, then good reports and bad reports will not be systematically revealed. If, for example, 60% of all good reports are revealed, but only 10% of all bad reports are revealed, then the observed feedback is biased towards the positive. When the observed good and bad reports do not correctly represent the correct distribution of positive and negative transactions in the market, future buyers cannot update their beliefs correctly. Therefore, buyers' propensities to buy depend only on the initial prior belief μ_0^i and price *P*. In essence, it is the same as the case where no information is available.

In the second case, if the reporting benefits from giving good reports and bad reports are symmetric, then good reports and bad reports will be systematically revealed. Future buyers can update their beliefs correctly, but at a slower rate than in the case where all buyers report.[22] It takes longer for buyers to learn a seller's type in this case than in the complete information case.

The problem of missing feedback reports caused by reporting costs can be overcome by the following simple mechanism. The market maker provides sellers an *option* to provide rebates (not necessarily in monetary form) contingent on buyers' leaving feedback, regardless of whether the feedback is positive or negative.

Proposition 3. *If the rebate $r \leq \alpha - f(\mu_0)P$ and $r < f(\mu_0)P - \beta$, there exists an equilibrium where good sellers always provide rebates and bad sellers mimic good sellers by providing rebates even after their bad types are identified through the reports and it is no longer profitable to mimic*

---

[22] It is easy to illustrate this with examples. Consider a case where it is commonly know that α=0.8, β=0.3. Suppose a good seller shipped products 8 out of the first 10 periods and the 11th buyer's prior belief is μ_0=0.5. If the seller's past reputation history is fully revealed, then the 11th buyer's updated belief of the seller's type being good is $\frac{0.5\times0.8^8(1-0.8)^2}{0.5\times0.8^8(1-0.8)^2+0.5\times0.3^8(1-0.3)^2} = 0.99$. However, if the seller's past positive and negative feedback are symmetrically revealed half of the time, then the 11th buyer's updated belief is only $\frac{0.5\times0.8^4(1-0.8)^1}{0.5\times0.8^4(1-0.8)^1+0.5\times0.3^4(1-0.3)^1} = 0.93$. Please note that a good seller may not necessarily ship the products 8 out of every 10 times.

good sellers, given the off-equilibrium path belief that anyone who chooses not to rebate must be a

bad seller. If the rebate r > α- f(μ₀)P, then both types of sellers will choose not to offer rebates, NRS.

The intuition is as follows: if good sellers choose the rebate option but bad sellers do not, then buyers can immediately identify sellers' types by observing who provides rebates. Therefore, bad sellers will mimic good sellers in the early periods. Since both types of sellers provide rebates to overcome buyers' reporting costs, buyers will report. We assume that all buyers report honestly if they decide to report. Thus, a bad seller's type will be revealed through feedback reports. In the long run, even if a bad seller's type is revealed, he may keep on choosing the rebates. The reason is that if a buyer identifies a bad seller, she would not buy from him. Therefore, even if a bad seller chooses the rebate, he does not need to pay for the rebate. This pooling equilibrium still holds when we combine moral hazard with the adverse selection model.

Proof of Proposition 3:

Proof.

First, let's examine the separating equilibrium where good sellers choose rebates (*RS*),

and bad sellers choose no rebates (*NRS*). 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 he is a good seller, and the good seller's payoff is $\alpha - r$. If a seller does not choose the rebate option, then he is considered a bad seller; the buyer's willingness to pay to a bad seller is β, so the seller's payoff is β. If the rebate is larger than the price difference between good and bad sellers, i.e., $r > \alpha - \beta$ then both good and bad sellers choose not to rebate (*NRS*). If the rebate is less than the price difference between good and bad sellers, i.e., $r \leq \alpha - \beta$, we need to check whether any sellers want to deviate from the separating equilibrium. A bad seller would get the higher payoff $\alpha - r$ instead of β if he pretends to be a good seller by choosing the rebate option. Thus, the separating equilibrium does not exist.

Nor does there exist another separating equilibrium, where good sellers choose not to rebate and bad sellers choose to rebate. The payoff to the good seller is α, the payoff to the bad seller is $\beta - r$, and

$\alpha > \beta$, so that the bad seller can have a higher payoff if he presents himself as a good seller by choosing a rebate. The bad sellers have incentives to deviate from this separating equilibrium. Thus, the separating equilibrium does not exist, either.

Since the separating equilibrium does not exist, we next examine the pool equilibrium. First, we examine the pooling equilibrium where both types of sellers choose to provide rebates, *RS*. The rebate *r* is chosen in a form that can systematically reveal good and bad reporting.[23] To simplify the analysis, we assume that *r* is greater than the net reporting cost (i.e., $C - B$), so it will reveal all reports.[24] In this case, buyers cannot know a seller's type just by observing the seller's choice about providing rebates. However, since both types of sellers provide rebates, all buyers will provide reports. A future buyer can, by using the information about a seller's previous history, update her beliefs on the seller's type. If buyer *i* does not report, her expected utility at period *t + 1* is

$$EU^i_{b,t+1} = \mu^i_t\alpha + \left(1 - \mu^i_t\right)\beta - P, \qquad (1)$$

while if she chooses to report, her expected utility is

$$EU^i_{b,t+1} = \mu^i_t\alpha + \left(1 - \mu^i_t\right)\beta - P - C + B + r. \qquad (2)$$

In this case, *r* is greater than the net reporting benefit, so a buyer has higher expected payoff if she chooses to report. The expected payoff for the seller at period *t+1* is

$$EU^i_{s,t+1}(Rs; \text{if buyer chooses GR or BR}) = \Pr\left(\mu^i_t \geq \frac{P-\beta}{\alpha-\beta}\right) P - r = f(\mu^i_t)P - r. \quad (3)$$

As the number of time periods *t* becomes infinite, according to the Weak Law of Large Numbers, ?? for the good seller equals 1 with very high probability, and the buyer's willingness to pay converges to α; μ for the bad seller equals 0 with very high probability, and the buyer's willingness to pay converges to β. If the expected payoff from providing rebates is higher than the expected payoff from not providing rebates in the long run, i.e., $\alpha - r \geq f(\mu_0)P$, then the good seller will choose to provide the rebate if he is patient. If the expected payoff from providing the rebate is higher

---

[23] The form of rebate depends on what causes no reports or no systematic reports. In this paper, we use monetary form for the rebates. See Li (2010a) for more discussion.

[24] If *r* is not greater for all buyers' net reporting cost, some buyers will report. As long as good and bad reports are systematically revealed, future buyers will still be able to update their beliefs in the right direction, and at a slower rate than in the case where all buyers report as shown in footnote 26. The general results still hold.

than the expected payoff from not providing at the beginning, i.e., $f(\mu_0)P - r > \beta$, then the bad seller wants to mimic the good seller. In the long run, even when a bad seller's type is revealed i.e., $f(\mu_t)P - r < \beta$, he may continue to choose the rebates. The reason is that if a buyer identifies a bad seller, the buyer would choose not to buy from the seller. Therefore, the bad seller does not need to pay for the rebate.[25] So as long as $r \leq \alpha - f(\mu_0)P$ and $r < f(\mu_0)P - \beta$, the patient good sellers will choose to give the rebate, and the patient bad sellers will also give the rebate. If $r > \alpha - f(\mu_0)P$, both types of sellers will choose no rebate, *NRS*.

Another pooling equilibrium is that both types of sellers choose not to rebate, (*NRS*), supported by the off-equilibrium path belief that anyone who chooses to rebate, *RS*, must be a bad seller. In this case, the seller's expected payoff is $f(\mu_0)P$ for every period.

If $r \leq \alpha - f(\mu_0)P$, this equilibrium does not exist if we use the intuition criteria. Since 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 any sellers who chooses to rebate are bad, is not feasible. If $r > \alpha - f(\mu_0)P$, then the pooling equilibrium in which both types of sellers choose no rebate, *NRS*, exists. a? ?

Appendix D: Proof of footnote 25

Proof.

For each history h_t of feedback history, buyers use Bayes' rule to update their beliefs about sellers' types when there are complete histories about the sellers,

$$\mu_t(\theta_G|h_t) = \mu_{t_{GR}+t_{BR}} = \frac{\mu_0^i \alpha^{t_{GR}}(1-\alpha)^{t_{BR}}}{\mu_0^i \alpha^{t_{GR}}(1-\alpha)^{t_{BR}}+(1-\mu_0^i)\beta^{t_{GR}}(1-\beta)^{t_{BR}}}.$$

We can rewrite this equation as the following:

$$\mu_t(\theta_G|h_t) = \frac{1}{1+\frac{1-\mu_0^i}{\mu_0^i}\left[\left(\frac{\beta}{\alpha}\right)^{\frac{t_{GR}}{t}}\left(\frac{1-\beta}{1-\alpha}\right)^{\frac{t_{BR}}{t}}\right]^t}.$$

If the seller is of type θ_G, then the Weak Law of Large Numbers implies that

---

[25] This result is different than the result in the auction market case as shown in Li (2010a) where a bad seller will no longer give a rebate after his type is revealed through reports.

$$Prob\left(\frac{\tilde{t}_{GR}}{t} \approx \alpha\right) \approx 1.$$

When t is large, where $\tilde{t}_{GR}$ is a random variable that records the number of good reports up to and including period t. If $\frac{\tilde{t}_{GR}}{t} \approx \alpha$, then

$$\left(\frac{\beta}{\alpha}\right)^{\frac{t_{GR}}{t}}\left(\frac{1-\beta}{1-\alpha}\right)^{\frac{t_{BR}}{t}} \approx \left(\frac{\beta}{\alpha}\right)^{\alpha}\left(\frac{1-\beta}{1-\alpha}\right)^{1-\alpha}.$$

Therefore

$$\ln\left[\left(\frac{\beta}{\alpha}\right)^{\alpha}\left(\frac{1-\beta}{1-\alpha}\right)^{1-\alpha}\right] = \alpha \ln\left(\frac{\beta}{\alpha}\right) + (1-\alpha)\ln\left(\frac{1-\beta}{1-\alpha}\right) < \ln(\beta + 1 - \beta) = 0,$$

and it follows that $\left(\frac{\beta}{\alpha}\right)^{\alpha}\left(\frac{1-\beta}{1-\alpha}\right)^{1-\alpha} < 1.$ Consequently, for large t, we conclude that

$$\frac{1}{1+\frac{1-\mu_0^i}{\mu_0^i}[\left(\frac{\beta}{\alpha}\right)^{\frac{t_{GR}}{t}}\left(\frac{1-\beta}{1-\alpha}\right)^{\frac{t_{BR}}{t}}]^t} \approx \frac{1}{1+\frac{1-\mu_0^i}{\mu_0^i}[\left(\frac{\beta}{\alpha}\right)^{\alpha}\left(\frac{1-\beta}{1-\alpha}\right)^{1-\alpha}]^t} \approx \frac{1}{1+0} = 1 .$$

with high probability. A similar argument established that $\mu_t(\theta_B|h_t) \approx 0$ with high probability when the seller is of type θ_B and t is large. ∎