

MPRA

Munich Personal RePEc Archive

Money Talks? An Experimental Study of Rebate in Reputation System Design

Li, Lingfang (Ivy) and Xiao, Erte

29 April 2010

Online at <https://mpra.ub.uni-muenchen.de/22401/>
MPRA Paper No. 22401, posted 30 Apr 2010 02:13 UTC

Money Talks? An Experimental Study of Rebate 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

April, 2010

Abstract:

Reputation systems that rely on feedback from traders are important institutions for helping sustain trust in markets, while feedback information is usually considered a public good. We apply both theoretical models and experiments to study how raters' feedback behavior responds to different reporting costs and how to improve market efficiency by introducing a pre-commitment device for sellers in reputation systems. In particular, the pre-commitment device we study here allows sellers to provide rebates to cover buyers' reporting costs before buyers make purchasing decisions. Using a buyer-seller trust game with a unilateral feedback scheme, we find that a buyer's propensity to leave feedback is more sensitive to reporting costs when the seller cooperates than when the seller defects. The seller's decision on whether to provide a rebate significantly affects the buyer's decision to leave feedback by compensating for the feedback costs. More importantly, the rebate decision has a significant impact on the buyer's purchasing decision via *signaling* the seller's cooperative type. The experimental results show that the rebate mechanism improves the market efficiency.

JEL: C91, D82, L86, H41.

Keywords: reputation, trust, feedback mechanism, asymmetric information, public good, experimental economics.

Acknowledgements: We would like to thank Daniel Houser, Luis Cabral, seminar participants at George Mason University, Rutgers University, Carnegie Mellon University, Peking University, Fudan University, China Central University of Finance and Economics, Zhejiang University, Shanghai University of Finance and Economics, Baidu Inc., Alibaba Research Institute, and ESA International conference 2009 for valuable comments on the earlier drafts. We are grateful for grants to the first author from the University of Louisville Research Inventive program and the 211 Project for Shanghai University of Finance and Economics (the 3rd phase), and to the second author from the National Science Foundation (SES-0961341) that supported this research.

I. Introduction

Reputation systems play an important role in sustaining trust in exchange environments where information is asymmetric. In many situations, such as online trading, building a reputation system often relies on voluntary feedback from involved parties. Such reputation systems, however, may not be reliable due to the cost of reporting. Previous literature has suggested various mechanisms for improving the reliability of reputation systems. Most of those suggestions, however, either require reporters to bear the reporting cost or require the market to provide incentives. In this paper, we propose and examine a rebate mechanism that gives sellers the option to cover the reporting cost for buyers in a listing price market. We show, both theoretically and experimentally, that this mechanism improves the reputation system and market efficiency without requiring buyers or markets to bear the reporting cost.

Previous research has pointed out that reputation systems can enhance cooperation when information is asymmetric¹. With the easy accessibility of the Internet, online feedback systems play an important role in building reputation and trust to facilitate the interaction between principals and agents. For example, online markets like eBay and Amazon provide feedback systems that allow traders to obtain reputation information for their counterparts or the products on sale. Data, however, suggests that only about 50% of buyers leave feedback after transactions on eBay. Furthermore, the existing feedback is biased toward positive due to missing negative feedback (see, Resnick and Zeckhauser, 2002, Cabral and Hortacsu, 2010, and Dellarocas and Wood, 2008).

¹ See, Bajari and Hortacsu, 2004; Bolton and Ockenfels, forthcoming; Bolton et al., 2005; Brandts and Figueras, 2003; Brown and Morgan, 2006; Dellarocas, 2003a; Diamond, 1989; Fehr et al., 2009; Friedman and Resnick, 2001; Grosskopf and Sarin, forthcoming; Houser and Wooders, 2006; Jin and Kato, 2006; Klein, 1997; Kreps and Wilson, 1982; Milgrom et al. 1990; Shapiro, 1983; and Resnick et al. 2006.

There are two possible explanations for the biased feedback system. First, there is a direct cost associated with reporting feedback. Participants must take time and effort to leave feedback. Once feedback is provided, it often becomes public information that everyone can access. It may not directly benefit those who leave feedback. Thus, the feedback system endures the public goods type of social dilemma problem (see, Bolton et al. 2004 and Avery et al. 1999). Second, the missing feedback may also be attributed to the indirect cost of reporting caused by the design of the feedback system itself (see, Klein et al. 2005, Dellarocas and Wood, 2008, Bolton et al., 2009 and Li, forthcoming). For example, if both a buyer and a seller can leave feedback for each other after a transaction, the buyer may hesitate to leave negative feedback. The reason is that a strategic seller who receives negative feedback may retaliate against the buyer by leaving negative feedback. This possibility of retaliation could lead to an inefficient feedback system (e.g., Masclet and Pénard, 2008; Bolton et al., 2009). Some online markets, such as eBay, have realized the “fear of retaliation” problem. eBay even implemented a policy to ban sellers from leaving negative feedback for buyers in May 2008. Nevertheless, this type of mandatory policy has made some sellers switch to other sites and has discouraged sellers from leaving any feedback to buyers².

Several papers have suggested various mechanisms to induce buyers to leave feedback and thus induce sellers to cooperate. For example, Ba et al.(2002) suggests a Trusted Third Party (TTP) mechanism to issue certificates to sellers and buyers. Dellarocas (2003b) proposes charging a listing fee contingent on a seller's announced expected quality and rewarding the seller based on both the announced quality and the

² See <http://money.cnn.com/2008/05/14/smbusiness/ebayfeedback:fsb/index:htm> (accessed on September 28, 2008.)

rating posted for that seller by the winning bidder for that listing. Both papers suggest monitoring systems to induce sellers to cooperate. Miller et al.(2005)and Jurca and Faltings (2007) propose well-designed truth-eliciting incentive schemes to induce buyers to report honestly. Nevertheless, 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 will still lack incentives to leave feedback. If the market provides such incentives, it may burden itself with a huge volume of transactions.

In this paper, we study a rebate mechanism that facilitates exchanges and improves market efficiency without requiring buyers or the market to bear the cost. Under this rebate mechanism, sellers have the option to cover the reporting cost for buyers, regardless of whether the feedback is positive or negative. The essence of the rebate mechanism is that it plays a dual role of *incentivizing buyers to leave feedback* and *providing an option for sellers to signal quality or effort to cooperate*. Thus, it not only improves the reputation system, but also—and perhaps even more importantly—induces sellers to cooperate through the additional signaling mechanism.

We examine this rebate mechanism both theoretically and experimentally in a listing price market with a unilateral feedback system where the price of the product is predetermined and fixed and only buyers can leave feedback to sellers.³ We first show that in theory, under this rebate mechanism, there is a pooling equilibrium where both good and bad sellers choose the rebate option, even though their true types are revealed through feedback. This mechanism can induce bad sellers to cooperate and promote

³ As a first step toward examining the effectiveness of the rebate mechanism, we choose a listing price market and a unilateral feedback system. This simple setting provides us clean evidence to draw inferences regarding the signaling value of the rebate mechanism. See Li (2010) for the theoretical work on how the rebate mechanism works in auction markets.

efficient trades. To test the hypothesis derived from the theory and provide empirical evidence for the effectiveness of the rebate mechanism, we conduct experiments based on a buyer-seller trust game (see, Bolton et al., 2004). In the game, a buyer first decides whether to buy a product from a seller. If the buyer decides to buy the product, the seller then decides whether to ship it. Our experiment consists of four treatments. In the baseline treatment, the computer automatically and truthfully records the feedback for each seller. This treatment can be taken as a perfect reputation system (see, Bolton and Ockenfels, forthcoming). In the other two treatments, the feedback is reported by buyers at their own expense. These two treatments vary only in the magnitude of the reporting cost. The fourth treatment is based on the high reporting cost treatment where we introduce the rebate mechanism.

We investigate: 1) how reporting costs and rebate affect reporting behavior, and in particular, whether the reporting cost affects reporting of positive feedback in the same way as reporting of negative feedback?⁴; 2) how reporting costs and rebate affect trading behavior; 3) whether buyers take the rebate offer as a signal of a seller's quality; and 4) whether the rebate mechanism improves market efficiency.

We find that buyers are less likely to report when the reporting cost is higher. Interestingly, our data suggest that both the reporting cost and the seller's rebate decision have a significant effect on the buyer's propensity to leave feedback when the seller cooperates, but not when the seller defects. More importantly, the rebate decision has a significant impact on buyers' purchasing decisions via signaling the seller's cooperative

⁴ Experimental economics literature shows that there exists an asymmetry in positive/negative reciprocity, e.g., Offerman, 2002; Keysar et al, 2008. Al-Ubaydli and Lee, 2009.

type. Consistent with the signal, sellers are significantly more likely to ship when they offer the rebate than when they do not. Our data show that under the rebate mechanism, market efficiency increases with the frequency of rebates received by the buyers.

The remainder of the paper is organized as follows: Section II introduces the theory framework of the rebate mechanism in the costly reputation system. Section III discusses the experiment design. Theory prediction is presented in Section IV. Section V reports the empirical results. Section VI concludes.

II. The Mechanism

In this paper, we study the rebate mechanism in a listing price market environment.

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 numbers, so a buyer will not meet the same seller again.⁵ If the buyer decides to buy, then the seller decides whether to ship the product. Suppose a buyer values the product 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) = 0; U_s (buyer doesn't buy) = 0; U_b (buyer buys, seller ships) = $V_b - P$; U_s (buyer buys, seller ships) = $P - V_s$; U_b (buyer buys, seller doesn't ship) = $-P$; U_s (buyer buys, seller doesn't ship) = P .

⁵ Resnick and Zeckhauser (2002) report that 89% of all buyer-seller pairs conducted just one transaction during the five-month period covered by the data set.

II. A. Adverse Selection Model

To understand the intuitions behind the rebate mechanism, we start from a simple set up of a pure adverse selection model with infinite periods. We then consider a more complicated model with both adverse selection and moral hazard with finite periods.

Suppose that in a *pure adverse selection model with infinite periods*, where sellers behave according to their types, there exist two types of sellers: 1) a good type θ_G who has a high probability α of shipping the product; and 2) a bad type θ_B , who has a low 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 the $V_s = 0$ for both types of sellers, and $V_b = 1$ for buyers. Assume that both buyers and 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. 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 $1 - \mu_0^i$. When there is no information about the seller's past history, buyer i 's expected payoff is

$$\begin{aligned} EU_{b0}^i &= \mu_0^i (\alpha(1 - P) + (1 - \alpha)(-P)) + (1 - \mu_0^i) (\beta(1 - P) + (1 - \beta)(-P)) \\ &= \mu_0^i \alpha + (1 - \mu_0^i) \beta - P. \end{aligned}$$

If buyer i 's willingness to pay, $\mu_0^i \alpha + (1 - \mu_0^i) \beta$, is higher than the listing price P , then buyer i will buy even when there is no reputation information. Let $f(\mu^i)$ 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 $Pr(sale) = Pr(buy) = Pr(\mu^i \geq \frac{P - \beta}{\alpha - \beta}) = f(\mu^i)$. Since buyers and sellers are randomly matched in the market, a seller's expected payoff is

$EU_{s,0} = E(f(\mu_0^i) \cdot P) = P \cdot E(f(\mu_0^i))$ when there is no reputation information. To simplify notations, we denote $E(f(\mu_0^i))$ as $f(\mu_0)$ so $EU_{s,0} = f(\mu_0)P$.

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

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

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 initial round, a buyer whose initial belief μ_0^i is less than $\frac{P-\beta}{\alpha-\beta}$ dose not buy.

In the second round, after the buyer observes sellers' reputations, the buyer updates his/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 less) 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 seller⁶. 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 $\mu_0^i < \frac{P-\beta}{\alpha-\beta}$ will buy. If the true proportion of good sellers in the market is less than

⁶ Proof is provided in the Appendix A.

$\frac{P-\beta}{\alpha-\beta}$, 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 was $\mu_0^i \geq \frac{P-\beta}{\alpha-\beta}$ 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 due to the fact that 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 $b \in [0, B]$.⁷ If the reporting cost C is higher than the maximum reporting benefit B for all buyers, then no buyer reports.

If $C < B$, we consider two cases. In the first case, if the reporting 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 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

⁷ To report feedback can be a way buyers reciprocate sellers, 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; de Quervain et al. 2004; Fehr and Gächter, 2000; Xiao and Houser, 2005).

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.⁸ 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 1. *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 until their bad types are identified through the reports and are no longer profitable to mimic good sellers, given the off-equilibrium path belief that*

⁸ It is easy to illustrate this with examples. Consider a case where it is commonly known that $\alpha=0.8$, $\beta=0.3$. Suppose a good seller shipped products 8 out of the first 10 periods and the 11th buyer's prior belief is $\mu_0=0.5$. 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 $\frac{0.5 \times 0.8^8 (1-0.8)^2}{0.5 \times 0.8^8 (1-0.8)^2 + 0.5 \times 0.3^8 (1-0.3)^2} = 0.99$. 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 $\frac{0.5 \times 0.8^4 (1-0.8)^1}{0.5 \times 0.8^4 (1-0.8)^1 + 0.5 \times 0.3^4 (1-0.3)^1} = 0.93$. Please note that a good seller may not necessarily ship the products 8 out every 10 times.

anyone who chooses not to rebate must be a bad seller. If the rebate $r > \alpha \cdot f(\mu_0)P$, then both types of sellers will choose not to offer rebate, NRs.

Proof: See appendix B.

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 cost, 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, a bad seller will not have an incentive to mimic good sellers by providing the rebate. This pooling equilibrium still holds when we combine moral hazard with the adverse selection model.

II. B. Mixed Model with both Adverse Selection and Moral Hazard

In a model with both adverse selection and moral hazard with a limited period T where sellers can make an effort to influence transaction outcomes, and using a set up similar to the pure adverse selection model discussed above, 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 product with probability 1.⁹ 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 seller needs to consider whether to make the effort as part of his strategy. A buyer's propensity to buy at period t is $f(\mu_{t-1}^i + (1 - \mu_{t-1}^i)\hat{e}_t)$, where \hat{e}_t is the buyer's expectation of a bad seller's

⁹ It is possible to assume this probability to be φ ($0 < \varphi < 1$), and the main result still holds.

effort. The belief that the seller is a good type in period 2 after seeing a good report in period 1 is:

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

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 willingness to pay him is 0 thereafter. In the last period, T , the buyer's willingness to pay is μ_T , since bad sellers will not make an effort in the last period.

Proposition 2: *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 C.

When the reporting cost is more than the maximum internal reporting benefit, buyers will not report. Thus, bad sellers will have no incentive to make an effort. In this case, if we introduce the rebate mechanism into the market, it will induce bad sellers to put forth effort if the cost of providing a rebate is lower than the future gains from having a good reputation. This intuition is similar to that in the case of pure adverse selection.

Proposition 3: *If the rebate $r < [1 - (1 - \delta)f(\mu_0)]P - e(1)$ and $e(1) \leq [1 - (1 - \delta)f(\mu_0)]P$, then a bad seller will choose the rebate option and make an effort in the first $T - 1$ periods, but not in the last period T .*

Proof: See Appendix D.

Unlike a monitoring mechanism suggested in the earlier literature (e.g., Ba and Pavlou, 2002 and Dellarocas, 2003b), this rebate mechanism provides an *option* for sellers to *signal* quality rather than *requiring* them to do so. As an “inducing” mechanism, rather than an “enforcing” mechanism, it allows both good sellers and bad sellers to co-exist in the market, but makes it possible for buyers to distinguish between them through feedback over time.

We next describe the experiments we conducted to examine empirically how people behave under this rebate mechanism.

III. Experiment

III. A Experiment Design

Our experiment consists of four treatments: 1) Computer automatic feedback treatment (*Auto_fb*); 2) Feedback cost 10 treatment (*C10*); 3) Feedback cost 5 treatment (*C5*); and 4) Feedback cost 10 with Rebate 5 treatment (*C10r5*). Each treatment consists of 10 rounds. In each round, each subject is randomly paired with another subject. 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 (Bolton and Ockenfels, 2007) described in Figure A. It is exactly the same in each treatment. 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 ends 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 earns 50E\$. If the seller does not ship the product, the seller earns 60E\$ and the buyer earns 10E\$.

The first stage is followed by a second stage where the seller receives feedback. The seller is informed of the feedback she 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. This treatment informs us what happens under the perfect reputation system.

Feedback cost 10 treatment (C10) and Feedback cost 5 treatment (C5). (see. Figure B and C)

In both treatments, 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 above, 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 non-monetary value of reporting. We speculate that a reporting cost of E\$5 is probably not high enough for some buyers and that their psychological benefit of reporting may be higher than the monetary cost of E\$5. Thus, some buyers will report in this treatment. Our goal is to test the 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 E\$10. This is the highest reporting cost we could impose so that the buyer would not have a negative earning from the game (when the seller did not ship the product and the buyer chose to report). 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 her feedback for that round. In addition, if the buyer does not purchase the product, the paired seller also receives a feedback of “N/A (no report)” in that round. That is, if the seller has a feedback of “N/A (no report),” this means that either the buyer did not purchase the product from the seller or the buyer purchased but did not report feedback. The reason we design the feedback mechanism this way is 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) (see Figure D)

To examine the rebate mechanism, we introduce a rebate mechanism in the C10 treatment where we assume the high cost of reporting leads to an ineffective reputation mechanism. The design of the C10r 5 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. 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 on the non-monetary value of sellers' rebate choices to buyers¹⁰.

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 below, in this treatment we do not observe any incidence of

¹⁰Note 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 Auto_fb treatment).

feedback decisions that are inconsistent with the seller's shipping decision (i.e., positive feedback for "not ship" and negative feedback for "ship")¹¹.

III. B Procedure

The experiment was conducted at P.E.E.L lab 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 paired with another subject each round. One round was randomly chosen as the payoff round. Each subject was paid according to the outcome in that round. Subjects were paid privately.

IV. Theory predictions

In this section, we derive theory predictions of sellers' and buyers' decisions based on our theory framework discussed in Section II. Both the pure adverse selection model and the mixed model may be applied to our environment.¹² If the theoretical predictions for both models are different, we list them separately.

Buyer's reporting decisions

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

¹¹ In all three treatments where buyers can leave feedback, we only observe one incidence of inconsistent feedback. In the C5 treatment where no rebate option is available at all, one buyer in the second round left a negative feedback to a seller who 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.

¹² Our data from the survey also suggest that this is the case. For example, one of the participating sellers wrote on the post-experiment survey that she had made up her mind to cheat in the 4th and 8th round before she played the game, so her behavior is consistent with the pure adverse selection model.

$$Prob(\text{report}): 0 \cong C10 < C5 < 1$$

This hypothesis is derived from the assumption that buyers are risk neutral and utility maximizers, and the reporting cost of 10E\$ is so high that the buyer will not report in this treatment.

Buyers' purchasing decisions

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

- *pure adverse selection*

$$Prob(\text{buy}): C10, C5 < \text{Auto_fb}$$

- *mixed model*

$$Prob(\text{buy}): C10 < C5 \cong \text{Auto_fb}$$

In the pure adverse selection model, sellers act according to their types. In the Auto_fb treatment, information is complete, so buyers can correctly update their prior beliefs to decide whether to buy. If a seller's past history contains a lot of good reports so that it makes both optimistic and pessimistic buyers updated beliefs $\mu_t^i \geq \frac{P-\beta}{\alpha-\beta}$, then both types of buyers will buy. In the C10 treatment, as we discussed above, we assume that no one will report. Buyers cannot correctly update their beliefs, so only optimistic buyers whose initial belief is $\mu_0^i \geq \frac{P-\beta}{\alpha-\beta}$ will buy. In the C5 treatment, some buyers may report if the information about positive and negative is systematically revealed. Then buyers can update their beliefs correctly and both optimistic and pessimistic buyers will buy from good sellers, but at a slower rate than in the Auto_fb treatment. In this case, the probability of buying in C10 is lower than in C5 and Auto_fb. If the information is not

systematically revealed, then buyers cannot correctly update their beliefs, so only optimistic buyers will buy. In this case, the probability of buying in C10 is the same as in C5 and lower than in Auto_fb.

In the mixed model, a good seller always makes an effort, so s/he will always get positive feedback, if any. A bad seller may choose to make an effort. If he/she does not make an effort in some rounds and the buyers in those rounds report, then he/she will definitely receive some negative feedback. As a consequence, no buyer will buy from her in the future, since 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 in this treatment. Consequently, bad sellers will not make an effort, and only optimistic buyers will buy. In the C5 treatment, as we mentioned above, we assume that for some buyers, the benefit of reporting is higher than the reporting cost of \$E5, and as a result, they will report. Bad sellers know that if they fail to make an effort and at least some buyers in some rounds report, they will receive negative feedback. Consequently, they will make an effort in order to have future payoffs. Assuming both good and bad sellers will make an effort and ship the products, buyers in the C5 treatment will be more likely to buy the products than those in the C10 treatment. In the Auto_fb

case, all reports are revealed, so bad sellers will make an effort, and buyers will be more likely to buy than in the C10 case¹³.

Hypothesis 3: Buyers take “no rebate” as a signal of a bad seller.

$$Prob^{C10}(\text{buy}) > Prob^{C10r5}(\text{buy} \mid \text{No Rebate})$$

$$Prob^{C10r5}(\text{buy} \mid \text{Rebate}) > Prob^{C10r5}(\text{buy} \mid \text{No Rebate})$$

In both theoretical models, there exists a pooling equilibrium that 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 the type of seller, but “not providing 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.

¹³ 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 C10 treatment is indifferent from that in the C5 treatment. In the mixed model case, bad sellers will 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 is indifferent from C5 and Auto_fb treatments.

Seller's rebate and shipping decisions

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

- *pure adverse selection model :*

$$Prob(ship): C10 = C5 = Auto_fb$$

- *mixed model:*

$$Prob(ship): C10 < C5 \cong Auto_fb$$

In the pure adverse selection model, since sellers act according to their types, the probability of shipping should not change in different reporting cost treatments.

In the mixed model, the analysis of sellers' shipping decisions is similar to the analysis in Hypothesis 3. The reason is that 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 and Auto_fb treatments, at least some buyers will report, so bad sellers will make an effort in order to have future payoffs. As a result, the probability of shipping is higher than in the C10 treatment¹⁴.

Hypothesis 5: Seller's decision not to offer a rebate signals his/her intention to defect.

$$Prob^{C10r5}(ship | Rebate) > Prob^{C10r5}(ship | No Rebate)$$

¹⁴ If our assumption that no one reports in the C10 treatment is invalid (i.e., some buyers still report or sellers believe that some buyers will report), then bad sellers will make an effort to ship the products like in the C5 and Auto_fb treatments. As a result, the probability of shipping in the C10 treatment is indifferent from the C5 and Auto_fb treatments.

Our theoretical framework discussed above suggests that bad sellers will mimic good sellers by choosing the rebate option until it is no longer 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.

Market efficiency

Hypothesis 6: Market efficiency can be improved by the rebate mechanism: a buyer who receives rebates more often can also achieve more efficient trades.

We define efficient trades as cases where the buyer buys and the seller ships. As discussed in Hypothesis 3, buyers are more likely to buy when offered a rebate. Meanwhile, Hypothesis 5 suggests that the shipping rate is higher when sellers offer a rebate than when sellers do not. Combining these two hypotheses, we predict that the more often a buyer receives rebates, the more efficient trades the buyer can achieve.

V. Results

We obtained observations on 214 subjects: 24 pairs in Auto_fb; 23 pairs in C5; 24 pairs in C10; and 36 pairs in C10r5. We first report results of buyers' decisions, followed by results of sellers' decisions. Lastly, we examine how the rebate mechanism affects market efficiency.

V. A. Buyer's feedback reporting decisions

Our data suggest that buyers report feedback honestly most of the time in the C5, C10 and C10r5 treatments. One buyer in the C5 treatment paid 5E\$ to leave feedback of “negative” in one round, and one buyer in the C10 treatment paid 10E\$ to leave feedback of “neutral” in one round after the seller shipped the product. Also, one buyer in the C10r5 treatment left feedback of “neutral” after her seller provided a rebate and shipped the product. In all the other cases where the buyer chose to leave feedback, he/she left positive feedback for cooperative sellers and negative feedback for non-cooperative 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. We then compare buyers’ reporting behavior in the C10r5 treatment with C5 and C10 to explore the effect of rebate on feedback reporting behavior. Our main findings in this regard are that both reporting cost and rebates have a significant effect on the buyer’s propensity to leave feedback when the seller shipped the product but not when the seller failed to ship the product.

***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 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 left feedback. Using the data from Treatment C10 and Treatment C5, we ran a random 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. The coefficients of the two treatment dummies are significantly different (chi-square test, $p=0.01$)¹⁵.

Previous research has shown that people are more willing to punish bad behavior than to reward good behavior (see Al-Ubaydli and Lee, 2009). If reporting feedback is a way for the buyer to reciprocate to the seller, we should expect that buyers might 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 view 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. The propensity of buyers to leave feedback is more sensitive to the reporting cost when the seller is cooperative than when the seller defects.

We find that 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 negative biased.

To provide statistical evidence, we ran a random effect Logit regression analysis of feedback reporting behavior (see Table 2 column (2)) similar to the regression analysis

¹⁵ All the coefficient comparison tests reported henceforth are chi-square tests with two-tail p-values.

reported in Table 2 column (1). The only difference is that we separated the cases of cooperative sellers from the non-cooperative 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 than that of C10noship and the coefficient of C5ship is significantly different than that of C5noship ($p < 0.01$).

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.

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

Result 3. *The rebate affects a buyer's propensity to leave feedback when the seller cooperates but has no effect on the buyer's reporting propensity when the seller defects.*

To see the effect of a rebate on buyers' reporting decisions, we expand the random effect Logit regression model we discussed 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 result of the regression is 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). Meanwhile, the coefficient of “C10r5_norebship” is significantly different from “C5ship” ($p=0.02$), but not significantly different from “C10ship” ($p=0.56$).

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, by compensating for the cost, can make buyers more likely to leave feedback when sellers ship the product, but have no effect on buyers’ reporting decisions when sellers fail to ship the product. This suggests that if buyers’ reporting behavior is negative-biased as shown above, then offering a rebate can reduce such negative bias by increasing the reporting rate when sellers cooperate. It is also worthwhile to note that, as we mentioned earlier, buyers report feedback honestly most of the time regardless of whether a rebate is provided. In other words, our data suggest that a seller’s rebate decision affects the propensity to report, but not the honesty of the feedback.

V. B. Buyer’s 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 about 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, as we report below, these differences are not significant. We start our analysis with a simple random effect Logit regression model of buyers' purchasing decisions, including only round, last round dummy, and three treatment dummies as independent variables. The result is 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.** Buyers' purchasing behavior is affected by the previous feedbacks sellers received.*

To gain further understanding about how buyers make the purchasing decision in each round, we next 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 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, Bolton et al., 2005). In view 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.

Therefore, 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 ($\sum_{t=1}^{t-2} \text{Positive fb}_{i,t}$), total number of negative feedback the seller received up to round t-2 ($\sum_{t=1}^{t-2} \text{Negative fb}_{i,t}$), whether the seller received positive feedback in the previous round ($\text{Positive fb}_{i,t-1}$), whether the seller received negative feedback in the previous round ($\text{Negative fb}_{i,t-1}$) and whether the seller shipped the product in the previous round ($\text{Ship}_{i,t-1}$). The result is reported in Table 3 column (2).

Consistent with previous literature (see the survey paper Dellarocas 2003a), the regression result shows 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 just 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 buy the product (the coefficient of $\sum_{t=1}^{t-2} \text{Negative fb}_{i,t}$ 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 decisions (the coefficient of $\sum_{t=1}^{t-2} \text{Positive fb}_{i,t}$ is positive but not significant).

We next investigate the effect of the rebate mechanism on buyers' purchasing decisions.

Result 5. *Supporting Hypothesis 3, buyers are significantly less likely to buy the product when the seller did not provide the rebate in 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 of the comparison of purchase rate in each condition, we further expand our regression model discussed above (Table 3 Column 2 regression) by including data from all four treatments. In addition to all the independent variables in the Column 2 regression, we add dummies of whether the seller provided a rebate to the buyer in this round ($C10r5_noreb_{i,t}$ and $C10r5_reb_{i,t}$). The result of the regression is reported in Table 3 column (3).

The regression result supports 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 significant lower than that of $C10r5_reb_{i,t}$, $p < 0.01$). Second, when the seller does not provide the rebate, his/her 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 $C5$, $p = 0.75$). These results suggest that *sellers’ rebate decisions affect buyers’ purchasing decisions not via compensating the ex post reporting cost but*

by signaling the buyer about the quality of the seller. We next provide evidence that the seller's rebate decision is indeed a credible signal of her cooperativeness.

V. C. Seller's 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 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 seller's shipping decision overall is significantly affected by the feedback reporting cost. The regression result is 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 suggests that the shipping rate is decreasing in the reporting cost, this effect of 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 6.** 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 1.

As we reported above, most buyers 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 his/her 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, she 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 significant ($p=0.15$), 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 significant ($p=0.01$).

Furthermore, we see more and more sellers offering the rebate over time. Figure 2 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 chance of making a profit.

Result 7. *The more often rebates a seller provides, the more likely the seller will ship the product.*

To examine whether sellers' shipping decisions are correlated with their rebate decisions, we calculate, for each seller, his/her 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 the 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 ($p < 0.01$).

V. D. Market efficiency

To test Hypothesis 6 that a rebate mechanism can improve market efficiency, 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 among each treatment.

***Result 8.** 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.*

Figure 3 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 3 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 $=0$ 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 10 rounds for each buyer:

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

To examine the effect of rebate on market efficiency, for each buyer in the C10r5 treatment, we calculated, over the 10 rounds, the average proportion of times she was offered a rebate before deciding whether to purchase the product:

$$\text{Avg_rebtms}_i = \begin{cases} \frac{\sum_{t=1}^{10} \text{C10r5_reb}_{i,t}}{10}, & \text{if C10r5 treatment} \\ 0, & \text{if Auto}_{fb}, \text{C5 or C10 treatment} \end{cases}$$

where $\text{C10r5_reb}_{i,t} = 1$ if buyer i was offered a rebate in round t in the C10 r5 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 result is 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 if a buyer always receives a rebate from the matched seller (i.e. $\text{Avg_rebtms}_i=1$), then the number of efficient trades can be as many as in the Auto_fb treatment.

Result 9. *When the rebate mechanism is available, buyers' and sellers' earnings are increasing in the number of rebates offered by sellers.*

Similar to the analysis of the effect of rebate on efficient trades, we examine the effect of rebate on the buyer and the seller's earnings using the same method. We calculate the average earnings over 10 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 result is shown in Table 6. We find the coefficients of C10 and C5 are significant 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 facilitates buyers to make efficient trades.

Using the same regression analysis, we next discuss the effect of reporting cost and the rebate mechanism on the sellers' earnings. The regression result is reported in Table 6. We find that 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 result also suggests 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, she 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 rebate every time, her earnings will be just slightly lower than in the Auto_fb treatment.

VI. Discussion

In this paper, we use theoretical and empirical methods to study an innovative rebate-for-feedback mechanism in reputation systems. Our experimental data support our hypothesis about the effectiveness of the rebate mechanism. In particular, the rebate offer from the seller provides buyers a credible signal as to the seller's quality. Sellers who offer more rebates are more likely to cooperate. Importantly, the rebate mechanism can improve market efficiency in that the more often rebates are offered, the more efficient trades occur. The market designer may thus consider incorporating the rebate option mechanism into the real market when there is an asymmetric information problem.

We also find that a buyer's propensity to report is more sensitive to reporting cost when a seller cooperates than when a seller defects. Consistent with this result, a rebate offer increases a buyer's propensity to leave feedback when the seller cooperates, but has no effect when the seller defects. In other words, the feedback we observe from our

experiment is negatively biased. This is in contrast to the empirical observations from eBay's website, where 99% of feedback was positive in the eBay-like bilateral feedback system (See Resnick and Zeckhauser, 2002, Cabral and Hortacsu, forthcoming).¹⁶

One explanation is that the bi-lateral feedback system, in contrast to the unilateral feedback reporting system studied in this paper, creates asymmetric costs for reporting (the cost for reporting negative feedback is higher than that for reporting positive feedback). As we mentioned above, one of the reasons for these asymmetric costs is the fear of retaliation. When the feedback system is bi-lateral, a strategic seller has the opportunity to leave negative feedback for a buyer who leaves him negative feedback.¹⁷ In our experiment, only buyers can leave feedback to sellers, not the other way around.

Another possible explanation is that negative feedback reporting, as a way to punish the counterpart, is determined by whether the reporter thinks the counterpart has a negative intention. In our experiment, the lower payoff is caused only by the seller's dishonesty. Thus, it is always clear to buyers whether the seller intentionally cheated. However, in the naturally occurring online trading market, the final trading outcomes always bear some degree of uncertainty. For example, the seller might have shipped the product on time, but the package was delayed by the post office. Xiao and Kunreuther (2009) find that people tend not to punish when they are not sure whether there were bad intentions behind a bad outcome. It follows that buyers might tend not to leave negative feedback if they are uncertain about whether a seller harmed them intentionally.

¹⁶ This data was collected prior to the eBay's policy to ban sellers from leaving negative feedback to buyers in May 2008.

¹⁷ Li (forthcoming) and Klein et al. (2005) suggest that fear of retaliation is a concern buyers have about leaving negative feedback.

In this paper, we take a first step to test the effectiveness of the rebate mechanism in a listing price market. The rebate mechanism takes the form of monetary compensation. In the naturally occurring online market, the forms of rebates may depend on the real reporting cost or concerns for users. This provides direction for further research on this topic. We are also conducting further studies to examine the rebate mechanism in other types of markets, such as an auction market where a rebate mechanism may benefit sellers by raising the price of the products.


References

- Andreoni, J., William Harbaugh and Lise Vesterlund** 2003. "The carrot or the stick: Rewards, punishments, and cooperation". *American Economic Review*, 93(3), 893–902.
- Avery, Christopher, Paul Resnick, and Richard Zeckhauser.** 1999. "The Market for Evaluations." *American Economics Review*, 89(3): 564-584.
- Al-Ubaydli, Omar, and Min Sok Lee.** 2009. "An experimental study of asymmetric reciprocity." *Journal of Economic Behavior & Organization*, 72(2): 738-749.
- 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.
- Bolton, Gary E., and Axel Ockenfels.** forthcoming. "The Limits of Trust in Economic Transactions - Investigations of Perfect Reputation Systems" In: *Trust and Reputation*, ed. Chris Snijders. New York: Russell Sage.
- Bolton, Gary E., Ben Greiner, and Axel Ockenfels.** 2009. "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., Elena Katok, and Axel Ockenfels.** 2005. "Cooperation among strangers with limited information about reputation." *Journal of Public Economics*, 89(8): 1457-1468.
- 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.
- 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-contingent 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.
- Diamond, Douglas W.** 1989. "Reputation Acquisition in Debt Markets." *Journal of Political Economy*, 97(4): 828-862.
- 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.** 2000. "Cooperation and Punishment in Public Goods Experiments." *American Economic Review*, 90(4): 980-994.
- 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.
- Grosskopf, Brit, and Rajiv Sarin.** forthcoming. "Is Reputation Good or Bad? An Experiment." *American Economic Review*.
- Houser, Daniel, and John Wooders.** 2006. "Reputation in auctions: Theory and evidence from eBay." *Journal of Economics and Management Strategy*, 15(2):353-369.
- 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.
- Klein, Daniel B.** 1997. *Reputation: Studies in the Voluntary Elicitation of Good Conduct*. Ann Arbor: The University of Michigan Press.

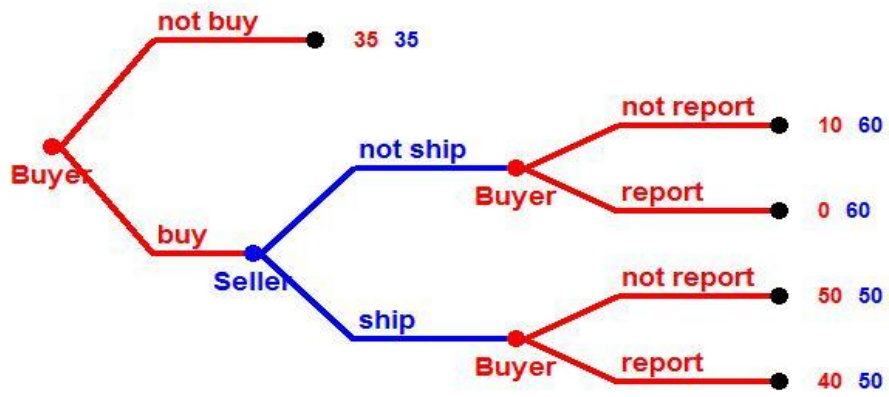
- Klein , Tobias J. , Christian Lambertz , Giancarlo Spagnolo , Konrad O. Stahl** and
 “Last minute feedback. ”SFB/TR 15. Discussion Paper (September 29, 2005), 2005.
- Kreps, David M., and Robert Wilson.** 1982. “Reputation and Imperfect Information.”
Journal of Economic Theory, 27(2): 253-279.
- Li, Lingfang (Ivy).** forthcoming. “What is the Cost of Venting? Evidence from eBay,”
Economics Letters.
- Li, Lingfang(Ivy).** 2010. “Reputation, Trust, and Rebates: How Online Auction Markets
 Can Improve Their Feedback Mechanisms.” *Journal of Economics and
 Management Strategy*, 19(2): 303-331.
- Masclot, 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.
- Milgrom, Paul R., Douglass C. North, and Barry R. Weingast.** 1990. “The Role of
 Institutions in the Revival of Trade: the Law Merchant, Private Judges, and the
 Champagne Fairs.” *Economics and Politics*, 2(1): 1-23
- 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.
- 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.
- Xiao, Erte , and Howard Kunreuther.** 2009 "Punishment and Cooperation in Stochastic
 Prisoner’s Dilemma Game". working paper.

Figure A. Seller-Buyer game (Auto_fb treatment)




Note: the first number is the buyer's payoff and the second number is the seller's payoff.

Figure B. C10 treatment



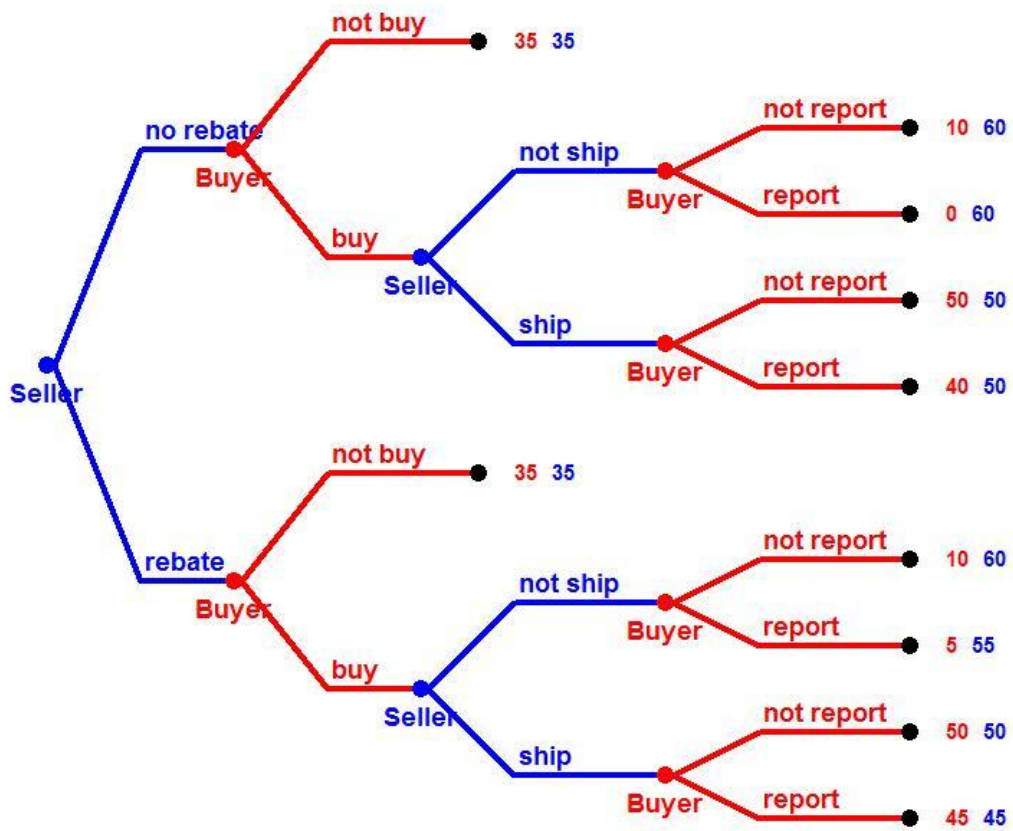
Note: the first number is the buyer's payoff and the second number is the seller's payoff.

Figure C. C5 treatment



Note: the first number is the buyer's payoff and the second number is the seller's payoff.

Figure D. C10r5 treatment



Note: the first number is the buyer's payoff and the second number is the seller's payoff.

Figure 1. Distribution of sellers' rebate frequency in the C10r5 treatment

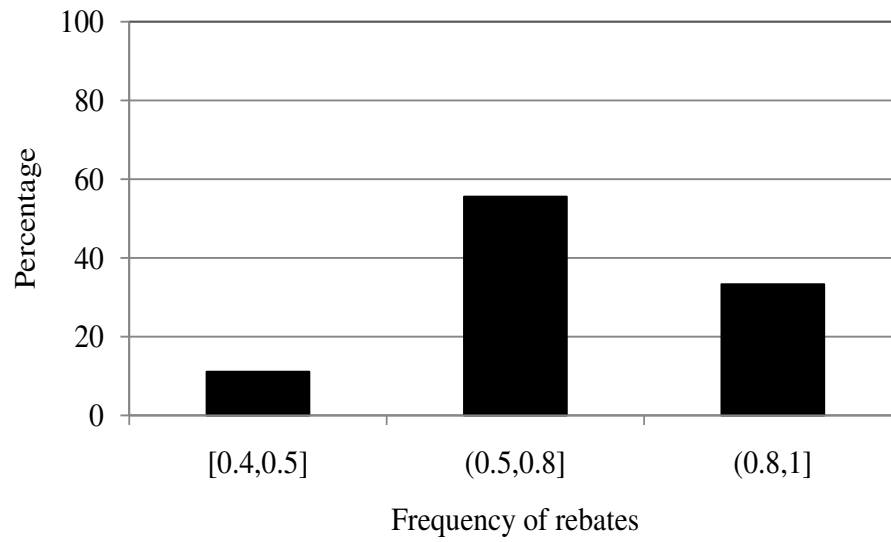


Figure 2. Proportion of sellers who offered rebates over 10 rounds




Figure 3. Proportion of efficient trades in each treatment




Table 1 Descriptive data of the decisions

Treatment	Buyer		Seller
	Purchasing rate	Reporting rate*	Shipping rate*
Auto_fb	79.17		88.42
C5	69.57	31.88	83.75
C10	67.50	10.49	81.48
C10r5_reb	83.91	27.85	85.84
C10r5_noreb	36.14	33.33	56.67

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

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

	Report _{i,t} (=1 if report in round t; =0 if not report)		
	(1) (include C5 and C10 treatments)	(2) (include C5 and C10 treatments)	(3) (include C5, C10 and C10r5 treatments)
	Coef. (s.e.)	Coef. (s.e.)	Coef. (s.e.)
Round(t)	-0.09 (0.07)	-0.44 (0.14)	-0.30 (0.08)
Final round (=1 if t=10)	-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

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

	Buy _{i,t} (=1 if purchased in round t; =0 if not)		
	Include Auto_fb, C5 & C10 treatments (1)	(2)	Include all the four treatments (3)
	Coef. (s.e.)	Coef. (s.e.)	Coef. (s.e.)
Round (t)	-0.04 (0.05)	0.07 (0.12)	-0.03 (0.09)
Final round (=1, if t=10)	-1.52 (0.42)	-2.74 (0.79)	-2.18 (0.56)
C10	1.77 (0.61)	2.37 (1.11)	2.77 (0.89)
C5	1.79 (0.60)	2.68 (1.14)	2.76 (0.89)
Auto_fb	2.37 (0.59)	0.92 (1.18)	1.08 (0.90)
$\sum_{t=1}^{t-2}$ Positive fb _{i,t}		0.02 (0.15)	0.12(0.13)
$\sum_{t=1}^{t-2}$ Negative fb _{i,t}		-2.48 (0.51)	-1.81(0.35)
Positive fb _{i,t-1}		1.77 (0.79)	1.56 (0.56)
Negative fb _{i,t-1}		-2.95 (0.88)	-2.75 (0.65)
Ship _{i,t-1}		1.61 (0.57)	1.31 (0.43)
C10r5_noreb _{i,t}			-1.79 (0.87)
C10r5_reb _{i,t}			2.51 (0.80)
Wald chi2(-)	45.75	50.84	89.56

Table 4. Sellers' Shipping Decisions: Regression Models

	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

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

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

Appendix A:

Proof of footnote 5:

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:

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

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

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

When t is large, where \tilde{t}_{GR} is a random variable (rv) 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

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

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. ■

Appendix B:

Proof of Proposition 1:

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

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 she is a good seller, and the good seller's payoff is $\alpha - r$. If a seller does not choose the rebate option, then she is considered as a bad seller; the buyer's willingness to pay is β , and 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 she 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 make good and bad reporting being systematically revealed.¹⁸ To simplify the analysis, we assume that r is greater than the net reporting cost (i.e., $C - B$), so it will make all reports being revealed.¹⁹ 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

¹⁸ The form of rebate depends on what causes no reports or no systematic reports. See Li (2010) for more discussion on this.

¹⁹ 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. The general results still hold.

$$EU_{b,t+1}^i = \mu_t^i \alpha + (1 - \mu_t^i) \beta - P, \quad (1)$$

while if she chooses to report, her expected utility is

$$EU_{b,t+1}^i = \mu_t^i \alpha + (1 - \mu_t^i) \beta - P - C + B + r. \quad (2)$$

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

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

As the number of time periods t becomes infinite, according to the Weak Law of Large Numbers, μ 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 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 good seller and choose the rebate until his payoff is less than one identified with a bad seller, i.e., $f(\mu_t)P - r < \beta$. 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 until their payoff $f(\mu_t)P - r$ is less than β , and choose no rebate otherwise. 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 anyone who chooses rebate is 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. ■

Appendix C:

Proof of proposition 2:

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 $EU_s = E(f(\mu_0^i) \cdot P) = P \cdot E(f(\mu_0^i))$.

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

$$EU_s = E[\Pr(\text{sale})_1 P - e(1) + \delta \Pr(\text{sale})_2 P] = E[f(\mu_0^i + (1 - \mu_0^i)\hat{e}_1)P - e(1) + \delta f(\frac{\mu_0^i}{\mu_0^i + (1 - \mu_0^i)\hat{e}_1})P],$$

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:

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

$$\text{If } e_1 = 1, I(e_1) = 1, \text{ and } V_1 = 1 + \delta V_2 - e(1).$$

$$\text{If } e_1 = 0, I(e_1) = 0, \text{ and } V_1 = E[\Pr(\text{sale})_1 P] = E[f(\mu_0^i + (1 - \mu_0^i)\hat{e}_1)P] = E(f(\mu_0^i)P) = f(\mu_0)P.$$

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

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

$$\text{If } e_2 = 0, I(e_2) = 0, \text{ and } V_2 = E[\Pr(\text{sale})_2 P] = E[f(\mu_1^i + (1 - \mu_1^i)\hat{e}_2)P] = E(f(\mu_1^i)P) = E(f(\mu_0^i)P) = f(\mu_0)P.$$

.....

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

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 $V_{T-1} = E[\Pr(\text{sale})_{T-1} P] = f(\mu_0)P$.

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

In order 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. ■

Appendix D:

Proof of Proposition 3:

Proof. If the reporting cost is more than the maximum reporting benefit, then no buyer will be inclined to report. In this case, the expected buyer's willingness to pay is $P_i = f(\mu_0)P$. 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. If we use the incentive mechanism proposed in the pure adverse selection model, both types of sellers will choose to give a rebate if the rebate is less than the price difference for good transactions and bad transactions, i.e., $r < E[\Pr(\text{sale})_{T-1} P + \delta I(e_{T-1})V_T - e_{T-1}(1)]$. In equilibrium $\Pr(\text{sale})_{T-1} = 1$, $I(e_{T-1}) = 1$, $V_T = f(\mu_0)P$, and $e_{T-1}(1) = e(1)$, the above condition is simplified to $r < [1 - (1 - \delta)f(\mu_0)]P - e(1)$. Also, bad sellers will put forth effort as long as their payoffs are more than $f(\mu_0)P$ for each period. ■