

# Remain Silent and Ye Shall Suffer: Seller Exploitation of Reticent Buyers in an Experimental Reputation System\*

Robert S. Gazzale<sup>†</sup>  
Tapan Khopkar<sup>‡</sup>

December 9, 2008

## Abstract

By providing incentives for sellers to act in a trustworthy manner, reputation mechanisms in many online environments can mitigate moral-hazard problems when particular buyers and sellers interact infrequently. However, these mechanisms rely on buyers sharing their private information about sellers with the community, and thus may suffer from too little feedback when its provision is costly. In this experimental study, we compare a standard feedback mechanism to one in which sellers can inspect a buyer's feedback-provision history, thus providing the buyer with incentives to share private information even when costly. We find fairly high trust and trustworthiness levels in all markets, with buyers showing a willingness to provide costly feedback, especially negative feedback, sufficient to induce seller trustworthiness. While we find, *ceteris paribus*, evidence that the availability of feedback-provision histories increases buyer trust by reducing missing feedback, it did not improve overall trustworthiness as this information enabled sellers to discriminate and act in a trustworthy manner less frequently with those who share information less frequently.

*JEL Classification:* C72, C73, C91, C92, D82, L14

*Key words:* experimental economics, trust, reputation, electronic markets

---

\*A previous draft of this paper was circulated as "Sharing Information with Reputation Systems: An Experimental Study." We thank Kan Takeuchi for programming, and Hyejin Rho for excellent research assistance. We thank seminar participants at the University of Pittsburgh and the ESA Annual Meeting for many helpful suggestions.. All errors are our own.

<sup>†</sup>Department of Economics, Williams College, rgazzale@williams.edu.

<sup>‡</sup>tapan.khopkar@gmail.com

## 1 INTRODUCTION

Consider a transaction in which a seller faces a moral-hazard problem—money already in hand but product not yet delivered, she has an opportunity to defraud the buyer, perhaps by not delivering the promised quality.<sup>1</sup> In many transactions, the threat of lost future custom may be sufficient to keep her trustworthy. However, in many interesting e-commerce environments, two key features may render such threats empty. First, while both the buyer and seller may interact repeatedly with the e-commerce platform, a particular buyer and seller may interact infrequently. Second, these transactions are often private-monitoring games—only the buyer knows whether he received the promised quality. Thus, the buyer’s threat to withhold future custom may well be insufficient, and, absent some mechanism for information sharing, other buyers will lack the information necessary to sanction untrustworthy sellers.

Previous work (not to mention eBay’s success) has shown that a well functioning reputation system can be quite effective in solving the moral-hazard problem and inducing trustworthy behavior in such environments.<sup>2</sup> Kandori (1992) shows that games with perfect private monitoring<sup>3</sup> may attain any mutually beneficial outcome, including trusting buyers and trustworthy sellers in this application, if each person carries a systematically revised label about his behavior (a reputation score). In theory, a reputation mechanism in an environment in which agents interact with particular agents infrequently but with the marketplace indefinitely can induce the same levels of trust and trustworthy behavior as a pair interacting indefinitely.<sup>4</sup> In the same way a buyer who repeatedly interacts with a particular seller can use the promise of future custom to entice a seller into providing high quality service, a reputation mechanism enables future buyers to condition behavior on a seller’s current actions. Hence, reputation systems can work as a feasible and less costly substitute for legal enforcement in electronic commerce (Bakos and Dellarocas 2003).

In Kandori’s model, information about an agent’s previous transactions becomes automatically available to her current partner. However, in environments with private monitoring, “feedback” (to use eBay’s terminology) must be voluntarily provided. This may be problematic in e-commerce environments as the incentives to provide feedback may be quite weak. Information sharing is at least marginally costly, and private benefits are likely small as the beneficiaries are likely strangers: the seller’s future potential customers.<sup>5</sup> This might

---

<sup>1</sup>To avoid ambiguity, we use feminine pronouns for sellers and masculine for buyers.

<sup>2</sup>In addition to providing incentives to act in a trustworthy manner (i.e., solving the moral hazard problem), reputation may also help in deciding who to trust (i.e., solving the adverse selection problem). In this study, we focus on the moral hazard problem. Dellarocas (2003) and Resnick, Zeckhauser, Friedman and Kuwabara (2000) provide an excellent overviews.

<sup>3</sup>Monitoring is perfect if an agent’s outcome perfectly reveals another’s action.

<sup>4</sup>Bolton, Katok and Ockenfels (2004) investigate this equivalence experimentally. They compare a market in which the same partners repeatedly trade to those in which “strangers” interact, either with or without a reputation mechanism. They find that while the reputation mechanism does improve trust and trustworthiness, the reputation market does not do as well as a partners market.

<sup>5</sup>When supply is constrained, there may be additional strategic motives for not sharing information. I might not want to tell you about my great babysitter if this puts my Saturday evening plans at risk.

well lead to an insufficient level of feedback.

At this point, we make a few observations about the “sufficient feedback” question. First, the level necessary to induce trustworthy behavior will depend on the environment. For example, with patient sellers and large gains from successful transactions, even a small feedback probability may be sufficient to deter opportunistic behavior.

Second, institutional features may encourage feedback provision. For example, eBay endeavors to make feedback provision relatively costless, and limiting feedback to three categories may reduce contemplation costs. Further, taking advantage of the fact that both sellers and buyers rate each other, many sellers on eBay encourage feedback with a policy of providing the buyer with feedback only after the buyer has left feedback.<sup>6</sup>

Finally, non-pecuniary incentives such as emotional satisfaction may be sufficient to induce some to provide feedback, particularly negative feedback.<sup>7</sup> Researchers have found in other environments a willingness to engage in costly punishment of those who violate trust. For example, Fehr and Gächter (2000) find in their experimental Voluntary Contribution Mechanism study community members willing to reduce the payments of a free rider even when costly to the punisher. Whereas contribution levels approached full free ride when a group had no opportunity to sanction, even groups whose composition changed each round achieved reasonably high contribution levels with the punishment option. In an experimental ultimatum game, Xiao and Houser (2005) find costly punishment more likely when subjects do lack other avenues to express their dissatisfaction toward low offers. This may be particularly relevant to e-commerce environments. Thus factors other than pecuniary incentives, such as joy of retaliation or altruism, induce sanctioning.

However, leaving feedback may not be as emotionally satisfying as monetary rewards or punishments. Whereas monetary punishments directly impose clear damage, undermining a seller’s reputation is an indirect and difficult to measure punishment. The damage, if any, takes place in the future, and may not be observable to the buyer. Furthermore, the punishment requires the participation of others (the seller’s future potential customers), and is thus uncertain. Thus satisfaction from providing feedback, and thus arguably incentives for providing this information, depends crucially on beliefs about the system’s effectiveness. For these reasons, the problem of inducing costly information provision in reputation systems invites further investigation.

There are a number of mechanisms available to induce information sharing. Direct payment for feedback by the platform operator is perhaps the most straightforward. Alternatively, noting that the trustworthy seller has greater incentives to elicit feedback, Li (2008) proposes a system in which the seller provides rebates to those who provide feedback.

An alternative solution is to allow buyers to develop reputations for information sharing by

---

<sup>6</sup>Of course, such a policy may reduce buyer incentives to leave negative feedback.

<sup>7</sup>We borrow eBay’s terminology and say that a buyer provides positive feedback if he reports that the seller shipped the item, and provides negative feedback if he reports that the seller did not ship the item.

giving the seller information about a buyer’s feedback-provision history. In much the same way the prospect of lost future custom helps “solve” the seller’s moral-hazard problem, a community denying the reticent buyer future gains from trade encourages him to leave feedback. Faced with a buyer with an information-sharing reputation, a seller, expecting to receive feedback, has an incentive to deliver the promised quality. Likewise, a buyer has incentives to provide feedback as developing a reputation for doing so results in better treatment from future sellers.

In order for such mechanism to work, the seller must be confident that denying the reticent buyer gains from trade—either by refusing to sell or not delivering the promised quality—will not result in negative feedback that harms *her* reputation. The availability of *second-order* histories can assist future buyers in discerning whether a seller’s failure to provide the promised quality was a community-sanctioned punishment of a reticent buyer or the cheating of a buyer in good standing.<sup>8</sup> Gazzale (2005) analyzes such a solution to the information-sharing problem in a repeated matching game in which players choose actions based only on own and partner reputations. Even if information provision is costly, such a mechanism yields a folk theorem-like result.<sup>9</sup> Full cooperation is attainable even when information sharing is costly.

However, the provision of second-order seller reputation information may be impractical. Its provision, implicitly sanctioning the punishment of reticent buyers, may be unpalatable to the trading platform. Further, second-order information greatly complicates the information processing required of market participants. More feasible is the provision first-order information: whether or not a seller provided the promised quality and whether or not the buyer provided feedback. In contrast with a mechanism requiring second-order reputation information, such a mechanism can be easily implemented in most practical reputation systems. In fact, eBay already makes easily available the feedback left by a member.

We thus consider a market in which both sides of a transaction have first-order reputation information about the other side. Buyers know whether or not the seller delivered the promised quality in previous transactions if feedback was left, but not the information-sharing reputation of the associated buyer. Sellers know whether or not the buyer provided feedback after previous transactions, but neither the content of the feedback<sup>10</sup> nor the reputation of the associated seller. Even in the absence of second-order information, sellers have the necessary information to differentiate among buyers. If sellers believe that a buyer’s feedback-giving history predicts the likelihood of feedback, then sellers can be selectively

---

<sup>8</sup>Alternatively, a central information aggregator can not report the negative feedback left by a buyer who had previously shirked on his information-sharing obligations.

<sup>9</sup>Any individually rational and feasible payoff can be supported as a subgame perfect equilibrium outcome of the indefinitely repeated game if players are sufficiently patient.

<sup>10</sup>We make this choice because the content of feedback, if visible, may complicate the buyers incentive to acquire a reputation for information sharing. If the buyer always leaves positive feedback, she might be labeled a pushover by the seller and the seller might attempt to defraud her. Likewise, if she always leaves negative feedback, the seller might believe fruitless any attempt to satisfy her. Thus, the buyer might seek the “optimal” ratio of positives to negatives in the spirit of Ely and Välimäki (2003), and not report truthfully.

trustworthy with some buyers. Buyer awareness of this discrimination gives him incentives provide more frequent feedback.

Even limiting available information to first-order information, a successful trading environment requires sophisticated strategic reasoning on the part of all participants. In the absence of the second-order information, discriminating based on a buyer’s feedback-giving history makes sense only if that history is a good predictor of the probability feedback. The seller must assess the risk that the previously reticent buyer leaves negative feedback. Likewise, the buyer seeking a reputation for information sharing must act proactively in developing this reputation. He must learn how frequently to provide feedback, and may need to learn, for example, that a strategy of only leaving feedback after untrustworthy behavior may be insufficient to develop this reputation.

In this study, we examine experimentally the ability of first-order reputation information to deliver the outcomes of trust, trustworthiness and costly feedback provision. We modify Bolton et al. (2004), allowing buyers to decide whether or not to leave feedback. Specifically, we study an environment in which a seller transacts with a randomly matched buyer in a trading game. The seller may act in a trustworthy manner or not, and the buyer can reveal, for a cost, this choice to the seller’s future partners. In one set of treatments (*NoHist*), the seller receives no information about the buyer’s information sharing history. This condition most resembles the original design of eBay’s Feedback Forum and many other reputation systems. In another set of treatments (*FbHist*), the seller automatically receives the buyer’s information-sharing record. In particular, a seller knows whether or not the buyer left feedback after all previous purchases, but not whether the feedback was positive or negative. We also vary across treatments the information sharing cost. Feedback profiles used in our experimental setup are similar to those used on most practical reputation systems.<sup>11</sup> Thus, the environment allows for buyer reputations and is simple enough to be implemented in practical reputation systems. Finally, we split each session into three subsessions, and “reset to zero” all reputation information between subsessions. We do this to study the development, if any, of trading game and information sharing norms.

In all treatments, we find relatively high levels of trust and trustworthiness supported by relatively high rates of negative feedback and relatively low rates of positive feedback. We find that feedback-provision histories are somewhat successful inducing buyer feedback. However, the availability of these histories does not improve trading game outcomes in this experiment, and actually leads to lower rates of trustworthy seller behavior in later subsessions. On the plus side, the reduction of missing feedback does promote additional trust for a given level of trustworthiness. However, sellers acting in a less trustworthy manner toward reticent buyers negates this positive effect.

---

<sup>11</sup>If the buyer does not give a feedback on a transaction, the feedback-giving history in our experiment reveals this fact. eBay, on the other hand, does not explicitly report that a transaction did not get a feedback. Thus it is not possible to precisely calculate the ratio of feedbacks to transactions on eBay.

## 2 EXPERIMENT DESIGN

In the beginning of each 12-subject session, the computer randomly and evenly assigns to subjects the role of either buyer or seller. Role assignments are fixed for the entire session. Each session consists of three subsessions each lasting 15 rounds. In each round, we randomly match buyers and sellers, and endow everyone with 35 points. The seller offers a hypothetical item for a sale at a price of 35. The buyer values the item at 50, and it costs the seller 20 to produce. The buyer first decides whether to buy from the seller and send the 35-point payment or not buy and keep his endowment. The round ends if the buyer chooses *not buy*. If the buyer decides to buy, then the seller decides whether or not to ship the item. If she ships the item, she pays the production cost and thus earns 50 ( $35+35-20$ ), whereas if she does not ship the item she does not pay the production cost and therefore earns 70 ( $35+35$ ). Thus, absent any information sharing costs, successful trade completion results in a 15-point surplus for both buyer and seller.

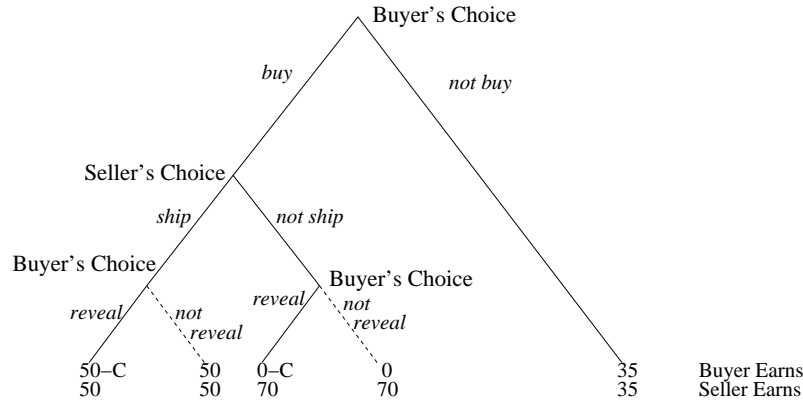


Figure 1: The game tree. The dashed line represents the action unavailable to buyers in the Baseline treatment.

If the buyer bought the item, he decides, after learning whether or not the seller shipped the item, whether or not to provide feedback, i.e., let this seller's subsequent partners know whether or not she shipped the item. He pays a cost  $C > 0$  to provide feedback, and incurs no cost if he chooses not to reveal the outcome. (If the buyer chooses *not buy*, the seller's future partners are automatically informed of this fact without cost to the buyer.) Note that a buyer only chooses whether or not to provide feedback. We ensure that that all feedback is truthful. It is not possible to report that the seller did not ship the item when in fact she did. Figure 1 depicts the game tree.

We employ a  $2 \times 2$  design to study the interactions between buyer information sharing histories and the sharing cost. In *FbHist* treatments, sellers have access to a buyer's feedback-leaving history. In *NoHist* treatments, a seller has no access to information about the buyer's past behavior, including whether or not he left feedback for previous sellers. In both the *FbHist* and *NoHist* treatments, we vary the feedback-provision cost (*High Cost*

vs. *Low Cost*). In the low-cost treatments, we set  $C = 2$ , whereas in the high-cost treatments we set  $C = 5$ .<sup>12</sup> In addition, we ran a baseline treatment (*Automatic*) in which feedback was automatically provided at no cost to the buyer.

<b>Information Condition</b>	<b>Cost Condition</b>		
	<i>No Cost</i> $C = 0$	<i>Low Cost</i> $C = 2$	<i>High Cost</i> $C = 5$
<i>FbHist</i> : Buyer feedback history available		<i>FbHistC2</i> 4 sessions	<i>FbHistC5</i> 4 sessions
<i>NoHist</i> : Buyer feedback history not available		<i>NoHistC2</i> 4 sessions	<i>NoHistC5</i> 4 sessions
<i>Automatic</i> : Feedback automatically provided	<i>Automatic</i> 4 sessions		

Table 1: Experimental treatments.

We conducted four sessions of each treatment in April and May 2006—two sessions at Williams College and two at the University of Michigan.<sup>13</sup> Table 1 details the treatments and sessions in this study. Participants were undergraduate students at each institution.<sup>14</sup> The experiment was programmed and conducted with z-Tree (Fischbacher 2007). Sessions lasted approximately one hour on average. The payoffs were in points, convertible at 100 points to 1 U.S. dollar. The average payoff was about \$24, including a \$3 show-up fee.

### 2.1 Information

In every round, the computer presents to both the buyer and seller the seller’s transaction summary and history. The transaction summary shows statistics about the seller’s past transactions in the current subsession: the number of items sold (i.e., the number of times previous buyers chose to buy from her), the number of positive and negative feedbacks received, and the number of rounds with no feedback. In addition to these summary statistics, the transaction history shows for each previous round in the current subsession whether the seller sold an item, whether feedback was left and, if feedback was indeed left, whether the seller did or did not ship the item.

In the *FbHist* treatments, the buyer also carries a history. Whereas a seller in the *NoHist* treatments has no information about her current buyer’s past actions, a seller in the *FbHist* treatments sees the buyer’s transaction summary and history. The buyer’s transaction summary shows statistics about his past transactions in this subsession: the number of items bought, the number of times he decided to leave feedback, and the number of rounds in which he did not provide feedback. The buyer’s transaction history shows, for each

<sup>12</sup>While these costs might be small relative to the item’s value, they are rather significant relative to the gains from trade: approximately 13% in the low-cost treatments, and 33% in the high-cost treatments.

<sup>13</sup>Baseline sessions were run in July 2007.

<sup>14</sup>Two Michigan students were graduate students.

previous round in the current subsession, whether the buyer bought an item and, if so, whether or not he provided feedback. *In no case, however, can a seller infer whether this previous feedback was positive or negative.* In all treatments, the computer displays to the buyer his own transaction summary and history.

In order to allow players the possibility of learning in this environment, we divide each 45-round session into three 15-round subsessions. Players retain the same roles, but at the end of the first and second subsessions, all transaction summaries and histories are reset. Thus, a buyer cannot access any information about a seller’s actions in a previous subsession, nor can a seller in a *FbHist* treatment access a buyer’s buying and feedback history from a previous subsession. We made the number of rounds in each subsession as well as the number of subsessions public knowledge in order to control for end-game effects.

## 2.2 Hypotheses

We use several performance measures to compare treatments. We measure *Trust* with the proportion of buyers who buy. We measure *Trustworthiness* with the proportion of items shipped conditional on being bought. We get a combined trust and trustworthiness measure by looking at *Efficiency*, defined as the proportion of potential transactions successfully completed. We also study information provision across treatments. We define the positive feedback rate (*PosRate*) as the proportion of *ship* decisions resulting in feedback, and the negative feedback rate (*NegRate*) as the proportion *not ship* decisions resulting in feedback.

If subjects can perform backward induction and believe others can as well, then, absent other-regarding preferences, in the unique subgame-perfect Nash equilibrium in any of our treatments sellers never ship and buyers never buy (or provide feedback). However, experimental results call into question the ability of subjects to perform the necessary backward induction (see, for example, Bolton, Katok and Ockenfels (2005), Johnson, Camerer, Sen and Rymon (2002), Selten and Stoecker (1986)). Furthermore, there is evidence for the existence of altruistic players in laboratory games (Andreoni and Miller 1993).

We therefore believe the appropriate model ought to include commitment sellers who ship and commitment buyers who provide feedback, despite explicit incentives to choose otherwise. In Appendix A, we identify the sequential equilibria of the treatments as we vary the fraction of sellers who are commitment sellers ( $\kappa_\sigma$ ) and the fraction of commitment buyers ( $\kappa_\beta$ ).<sup>15</sup>

When information-sharing histories are not available (the *NoHist* treatments), strategic sellers ship in the sequential equilibrium as long as the fraction of commitment buyers who always provide feedback is sufficiently high ( $\kappa_\beta \geq \frac{4}{7}$ ). If the fraction of commitment sellers

---

<sup>15</sup>In our theoretical model, we assume that each buyer and seller are only matched once. While each buyer/seller pair met 2.5 times per subsession on average, we do note that no subject who participated in a session in which we administered a post-experiment questionnaire reported attempting to infer previous interactions from published histories.



is sufficiently low ( $\kappa_\sigma < \frac{7}{10}$ ), the sequential equilibrium has the usual “chain-store paradox” flavor: at some round, strategic sellers start randomizing between *ship* and *not ship* and buyers between *buy* and *not buy* (with randomization occurring earlier for smaller values of  $\kappa_\sigma$ ). If the fraction of commitment feedback is not sufficiently high ( $\kappa_\beta < \frac{4}{7}$ ), then strategic sellers never ship and buyers never buy.

When information-sharing histories are available (the *FBHist* treatments), the minimum level of commitment feedback necessary for an equilibrium with strategic shipping is  $\kappa_\beta \geq \frac{8}{7\sqrt{7}}$ . Therefore, a lower level of commitment feedback is necessary when these histories are available. The nature of the sequential equilibrium also depends on the level of commitment shipping. When below a certain threshold,<sup>16</sup> there is no strategic shipping in the sequential equilibrium. When above a certain threshold,<sup>17</sup> strategic shippers choose *ship* in all but the final round. Thus, the level of commitment shipping required for an “always ship” sequential equilibrium is lower when information-sharing histories are available. Finally, when the level of commitment shipping is between the thresholds (and the level of commitment feedback is sufficiently high as defined above), commitment sellers ship and buyers buy and provide feedback for the first 12 rounds. In the 13th round, buyers buy, strategic sellers randomize, and strategic buyers randomize between providing feedback and not.

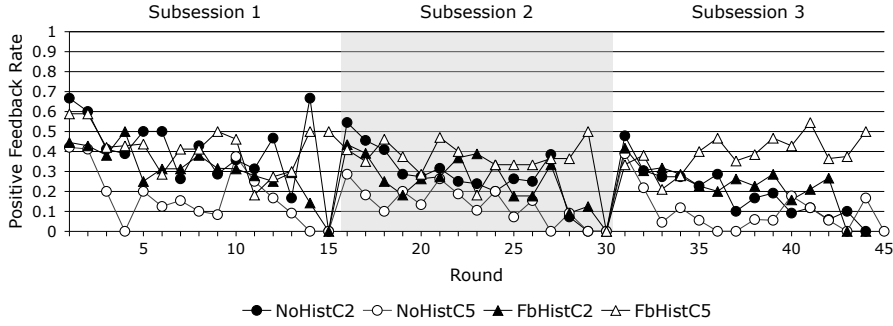
Therefore, while the availability of information-sharing histories introduces a requirement on the minimum level of commitment shipping necessary for a “good” sequential equilibrium, these histories enable good equilibrium outcomes for lower levels of commitment feedback. Coupled with a lower level of commitment shipping necessary for an “always ship, always buy” equilibrium, we hypothesize that the availability of feedback-provision histories will improve outcomes. We therefore make the following hypotheses.

**Hypothesis 1.** *Feedback rates will be higher in environments in which buyer feedback-provision histories are available to sellers.*

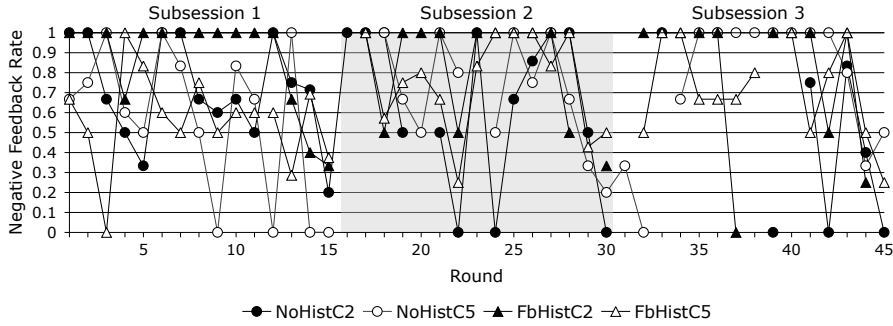
**Hypothesis 2.** *Trust will be higher in environments in which buyer feedback-provision histories are available to sellers.*

**Hypothesis 3.** *Trustworthiness will be higher in environments in which buyer feedback-provision histories are available to sellers.*

**Hypothesis 4.** *Efficiency will be higher in environments in which buyer feedback-provision histories are available to sellers.*



(a) Feedback Rate Conditional on Seller Choosing *ship*



(b) Feedback Rate Conditional on Seller Choosing *not ship*

Figure 2: Per-round feedback provision across treatments.

### 3 RESULTS

#### 3.1 Treatment Effects on Feedback

In all treatments, subjects demonstrated a willingness to leave feedback, particularly negative feedback. In Figures 2(a)–2(b), we depict feedback provision by round, aggregating for each treatment across all four sessions. A few observations become evident from the graphs. First, in all treatments positive-feedback provision starts out reasonably high, generally decreases somewhat over the a sub-session’s first few rounds, but then remains relatively constant until decreasing in the final rounds. Second, with the exception of the treatment in which feedback costs were high and histories were available (*FbHistC5*), the overall rate of positive feedback provision decreases from the first sub-session to the third. Third, in all treatments the provision of negative feedback starts out quite high and remains quite high (with a fair degree of inter-round variability) until dropping dramatically in the

<sup>16</sup>The exact threshold depends on the cost of feedback and the level of commitment feedback: it ranges from 0.015 to 0.034 when  $C = 2$  and from 0.038 to 0.084 when  $C = 5$ .

<sup>17</sup>The exact threshold depends on the cost of feedback and the level of commitment feedback: it ranges from 0.172 to 0.301 when  $C = 2$  and from 0.190 to 0.333 when  $C = 5$ .

Panel A:		Proportion feedback after ship in rounds 1–12				Permutation Tests	
Treatment	Sess. 1	Sess. 2	Sess. 3	Sess. 4	Overall	$H_1$	$p$ -value
<i>NoHistC2</i>	0.280	0.329	0.287	0.391	0.319	<i>NoHistC2 &lt; FbHistC2</i>	0.605
<i>FbHistC2</i>	0.435	0.246	0.337	0.215	0.305	<i>FbHistC5 &lt; FbHistC2</i>	0.739
<i>NoHistC5</i>	0.200	0.108	0.231	0.102	0.156	<i>NoHistC5 &lt; NoHistC2</i>	0.016**
<i>FbHistC5</i>	0.571	0.133	0.253	0.591	0.389	<i>NoHistC5 &lt; FbHistC5</i>	0.059*
Panel B:		Proportion feedback after not ship in rounds 1–12				Permutation Tests	
Treatment	Sess. 1	Sess. 2	Sess. 3	Sess. 4	Overall	$H_1$	$p$ -value
<i>NoHistC2</i>	0.875	0.654	0.550	0.846	0.687	<i>NoHistC2 &lt; FbHistC2</i>	0.040**
<i>FbHistC2</i>	0.913	0.813	0.944	1.000	0.908	<i>FbHistC5 &lt; FbHistC2</i>	0.014**
<i>NoHistC5</i>	0.622	0.864	0.875	0.636	0.732	<i>NoHistC5 &lt; NoHistC2</i>	0.609
<i>FbHistC5</i>	0.731	0.724	0.806	0.545	0.702	<i>NoHistC5 &lt; FbHistC5</i>	0.678

Note: Significant at: \* 10-percent level; \*\* 5-percent level; \*\*\* 1-percent level.

Table 2: Feedback provision by session over rounds 1–12 of all subsessions.

final rounds.<sup>18</sup> Finally, buyers appear more willing to incur negative feedback’s (direct) cost than positive feedback’s.

We formally analyze trading game performance across treatments by comparing outcomes over the first 12 rounds of all subsessions. We drop the last three rounds in order to control for end game effects. Panel A of Table 2 reports for each session the proportion *ship* decisions resulting in buyer feedback, as well as the alternative hypotheses and the  $p$ -values of the corresponding permutation tests with session as the unit of observation.<sup>19</sup> Panel B reports the same for negative-feedback. We now formally present results regarding treatment effects on feedback provision.

**Result 1** (Result: Feedback Provision.). *The availability of buyer feedback-provision histories weakly increases the provision of positive feedback when the cost of information sharing is relatively high, and significantly increases the provision of negative feedback when the cost is relatively low.*

**SUPPORT:** We report in the last two columns of Table 2 the corresponding alternative hypotheses and permutation test results. ■

We find limited support for the hypothesis that provision of buyer feedback-provision histories to sellers increases feedback provision (Hypothesis 1). These histories increased positive feedback only under relatively high feedback costs, and increased the provision of negative feedback only under the relatively low feedback costs.

<sup>18</sup>The small number of opportunities to leave negative feedback largely drives this variability. For example, in subsession three’s first 12 rounds in the four *FbHistC2* sessions, there were only 19 opportunities to leave negative feedback in the third subsession.

<sup>19</sup>The permutation test is a nonparametric version of a difference of two means t-test (Siegel and Castellan 1988, pp. 95–100). Pooling all independent observations, the  $p$ -value is the probability of observing a separation between the two treatments as the one observed when the pooled observations are randomly divided into two groups.

The relatively low rate of positive feedback in the treatment with high feedback costs and without feedback-leaving histories (*NoHistC5*) drives the positive feedback results. Feedback-provision histories increasing positive feedback only under relatively high feedback costs is consistent with a signaling story. Apparently, with low feedback costs, buyers felt little need to signal a willingness to provide negative feedback. However, with higher costs, the results are consistent with buyers taking advantage of the opportunity to signal a willingness to provide feedback. Whereas providing a signaling opportunity increases feedback provision’s benefits when the costs are high, directly decreasing the costs has a similar effect.

Likewise, the high negative-feedback rate in sessions with low feedback costs and available histories (*FbHistC2*) seems to drive the negative feedback results. It is perhaps helpful to think of these results in terms of complementarities. Taking away the signaling opportunity and increasing the feedback’s cost only have an effect on positive feedback when implemented together. Similarly, decreasing feedback’s cost and providing a means to signal only have an effect on negative feedback when implemented together.

**Result 2** (Result: Positive vs. Negative Feedback Rates). *The feedback rate is significantly higher after sellers choose not ship than after ship. That is  $NegRate > PosRate$  for all treatments.*

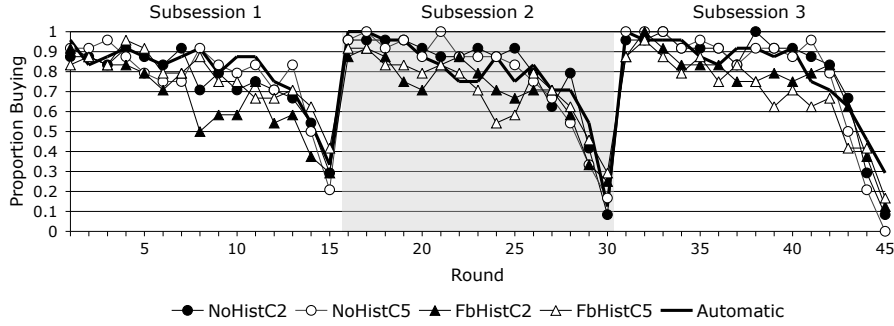
**SUPPORT:** For all treatments, the  $p$ -value (unreported) of the permutation test is  $p < .005$ . ■

The formal analysis confirms tendency toward the “always leave negative feedback, sometimes leave positive feedback” norm evident in Figures 2(a)–2(b). Of course, such a norm cannot be part of a subgame-perfect Nash equilibrium if with positive feedback costs net any non-pecuniary benefits. As an item’s non-receipt is private information, any buyer has a strict incentive to act as if he received the item and not leave costly feedback.


### 3.2 Treatment Effects on Trading Game Performance

In all treatments, levels of trust and trustworthy behavior, and thus efficiency, were quite high. In Figures 3(a)–3(c), we depict trading game outcome measures by round, where for each treatment we aggregate across all four sessions. A few observations become evident from the graphs. First, in all treatments, trust and trustworthiness (and thus efficiency) are quite high in a subsession’s first round, then gently decrease in subsequent rounds until dropping precipitously in a subsession’s final few rounds. Second, there appears to be a general improvement in trading game outcomes from the first to the third subsession. Third, low-feedback-cost treatments perform at least as well as the automatic-feedback baseline. Finally, the high-cost-of-information-sharing treatments have distinctly less trust, and somewhat less trustworthiness, than the low-cost-of-information-sharing treatments.


We formally analyze trading game performance across treatments by comparing outcomes



(a) Trust: Proportion buying.



(b) Trustworthiness: Proportion shipping conditional on selling.



(c) Efficiency: Proportion transactions successful.

Figure 3: Per-round trading game outcomes across treatments.

over the first 12 rounds of all subsessions, dropping the last three rounds to control for end-game effects. Panel A of Table 3 reports for each session the proportion of rounds in which the buyer bought, as well as the alternative hypotheses and the  $p$ -values of the corresponding permutation tests with session as the unit of observation. Panel B reports the same for shipping decisions conditional upon purchase and Panel C for the proportion of all potential transactions that result in the seller shipping the item. We now formally test our hypotheses regarding trading game outcomes.

**Result 3** (Result: Trading Game Outcomes.). *The availability of buyer feedback leaving*

Panel A:		Proportion buying in rounds 1–12				Permutation Tests	
Treatment	Sess. 1	Sess. 2	Sess. 3	Sess. 4	Overall	$H_1$	$p$ -value
<i>NoHistC2</i>	0.931	0.824	0.898	0.843	0.874	<i>NoHistC2</i> < <i>FbHistC2</i>	0.408
<i>NoHistC5</i>	0.657	0.829	0.815	0.819	0.780	<i>NoHistC5</i> < <i>NoHistC2</i>	0.032**
<i>FbHistC2</i>	0.852	0.866	0.907	0.898	0.881	<i>FbHistC5</i> < <i>FbHistC2</i>	0.070*
<i>FbHistC5</i>	0.833	0.898	0.625	0.787	0.786	<i>NoHistC5</i> < <i>FbHistC5</i>	0.454
<i>Automatic</i>	0.903	0.810	0.940	0.829	0.870	<i>FbHistC2</i> < <i>Automatic</i>	0.661
Panel B:		Proportion shipping in rounds 1–12				Permutation Tests	
Treatment	Sess. 1	Sess. 2	Sess. 3	Sess. 4	Overall	$H_1$	$p$ -value
<i>NoHistC2</i>	0.960	0.854	0.897	0.929	0.911	<i>NoHistC2</i> < <i>FbHistC2</i>	0.431
<i>NoHistC5</i>	0.739	0.877	0.909	0.938	0.872	<i>NoHistC5</i> < <i>NoHistC2</i>	0.262
<i>FbHistC2</i>	0.875	0.914	0.908	0.959	0.915	<i>FbHistC5</i> < <i>FbHistC2</i>	0.016**
<i>FbHistC5</i>	0.856	0.851	0.733	0.806	0.817	<i>FbHistC5</i> < <i>NoHistC5</i>	0.155
<i>Automatic</i>	0.918	0.843	0.956	0.899	0.904	<i>FbHistC2</i> < <i>Automatic</i>	0.636
Panel C:		Proportion successful transactions in rounds 1–12				Permutation Tests	
Treatment	Sess. 1	Sess. 2	Sess. 3	Sess. 4	Overall	$H_1$	$p$ -value
<i>NoHistC2</i>	0.894	0.704	0.806	0.782	0.874	<i>NoHistC2</i> < <i>FbHistC2</i>	0.429
<i>NoHistC5</i>	0.486	0.727	0.741	0.769	0.780	<i>NoHistC5</i> < <i>NoHistC2</i>	0.100*
<i>FbHistC2</i>	0.745	0.792	0.824	0.861	0.881	<i>FbHistC5</i> < <i>FbHistC2</i>	0.029**
<i>FbHistC5</i>	0.713	0.764	0.458	0.634	0.786	<i>NoHistC5</i> < <i>FbHistC5</i>	0.271
<i>Automatic</i>	0.829	0.676	0.898	0.745	0.870	<i>FbHistC2</i> < <i>Automatic</i>	0.629

Note: Significant at: \* 10-percent level; \*\* 5-percent level; \*\*\* 1-percent level.

Table 3: Outcomes by session over rounds 1–12 of all subsessions.

*histories increases neither trust, trustworthiness, nor efficiency.*

**SUPPORT:** We report in the last two columns of Panels A–C in Table 3 the corresponding alternative hypotheses and permutation test results. ■

Our results do not support the hypotheses that providing buyers’ feedback-provision histories to sellers improves trust (Hypothesis 2), trustworthiness (Hypothesis 3), and trading game efficiency (Hypothesis 4). We do note that as expected, trading game outcomes are generally deteriorating in information-sharing costs.

Aggregate trading game outcomes are somewhat surprising. First, while an increase in feedback costs decreased trust, we did not always observe a corresponding decrease in trustworthy behavior on the sellers’ part.<sup>20</sup> Second, feedback-provision histories did not improve trading game outcomes. Part of the explanation lies in the fact that there is little room for improvement—sessions without feedback-provision histories achieved rather high levels efficiency. Even without these histories, it appears as though sellers believed the net feedback cost, particularly for negative feedback, to be sufficiently low to make the threat of feedback credible. Sellers acted in a trustworthy manner to either develop a positive reputation or avoid a negative one. Relatedly, buyers, believing feedback’s threat credible, expected trustworthy behavior and were thus trusting.

<sup>20</sup>The relationship might work in both directions. More trust, particularly of sellers with a good reputation, increases the value of trustworthiness. Likewise, more trustworthiness ought to increase the value of trusting.

However, not only did information-sharing histories not improve trading-game outcomes, a case can be made that its availability actually worsened outcomes. Although not statistically significant, feedback-provision histories decreased trustworthiness when coupled with high feedback costs. Furthermore, while performance over all periods is important in assessing the desirability of an environment, it is also instructive to look at the play of the experienced. Dropping the first subsession, we weakly reject the null hypothesis of no difference in trustworthiness in  $C = 5$  sessions in favor of the alternate hypothesis of feedback-provision histories *decreasing* trustworthiness ( $p = 0.071$ ). Likewise, looking only at the final subsession, the decrease in trustworthiness arising from availability of these histories is weakly significant for both  $C=2$  ( $p = 0.100$ ) and  $C=5$  ( $p = 0.057$ ).

### 3.3 Treatment effects and trading dynamics

In this section, we investigate the underlying causes of the failure of feedback-provision histories to improve trading game outcomes despite the increase in feedback. We estimate a series of probit models to understand why buyers trust and sellers ship. In order to control for learning, we drop observations from the first subsession in all model specifications, and to control for end game effects we drop a subsession’s final three rounds. Further, we do not include observations from the baseline sessions in which feedback was automatically provided, as the information available to buyers was so markedly different.<sup>21</sup> In all models, we control for round, for subsession by including an indicator variable for subsession 3 ( $D_{subsess=3}$ ), and for high-feedback-cost sessions by including an indicator variable ( $D_{C=5}$ ) equal to 1 for those observations.

#### 3.3.1 Trustworthiness: The Ship Decision

In this section, we identify factors contributing to trustworthy behavior, with a particular emphasis on understanding why providing sellers with buyers’ feedback-provision histories did not improve outcomes. By looking at the factors that affected seller shipping decisions, we gain insights into the link between information available, and in particular information about a buyer’s feedback-provision history, and the shipping decision.

We find support for the hypothesis that the availability of feedback-provision histories does not improve trading game outcomes because sellers use this information to discriminate, acting in a trustworthy manner towards those with a reputation for leaving feedback, and on occasion taking advantage of those without. In building support for this hypothesis, we first look at the factors a seller may take into account even if she does not have access to the buyer’s feedback-provision history. While these factors contribute to her shipping decision, there remains a difference between treatments in which these histories were available and those in which they were not (i.e., the coefficient on the indicator variable for feedback-provision histories remains significant). We then look at only treatments with

---

<sup>21</sup>Inclusion of observations from *Automatic* sessions in models that do not include feedback-related regressors does not, however, qualitatively alter reported results.

feedback-provision histories, and find that differences in feedback-provision histories lead to differences in the probability that a seller ships the item.

Broadly speaking, two strategic reasons determine whether a seller chooses *not ship*.<sup>22</sup> First, she may choose *not ship* if she believes it will not be made known to future buyers. However, she may choose *not ship* despite a high probability of feedback if she believes that a bad reputation will not be so detrimental. While this will be the case if the future is short, it will also be the case if the probability of future purchase does not greatly depend on her reputation. Thus, holding constant feedback rates and the even probability of a sale with a “bad” reputation, more pessimistic beliefs about the probability of a sale with a good reputation will make *not ship* more enticing.

In order to assess the relative benefits of a good reputation, we first summarize the information a buyer has about a seller. As subjects did not have an exogenously given definition of a good reputation, we use the following definition. We say a seller has a good reputation in a particular round ( $goodRep = 1$ ) if the last feedback received was positive.<sup>23</sup>

The value of good reputation likely decreases in the number of rounds remaining. Additionally, a seller’s past returns from a good reputation may plausibly inform her assessment of its future returns. We thus define  $pctGoodNoSale$  as the ratio of the number of times in the current subsession a seller did not make a sale despite her good reputation to the number of rounds in which that seller had a good reputation. Intuitively, we expect that the fewer sales made despite a good reputation, the less likely she will ship. In order to distinguish between overall and more recent history, we define the indicator variable  $lastGoodNoSale$  equal to 1 if a seller did not make a sale in the most recent period in which she had a good reputation, and 0 if she did.

Holding constant the value of a good reputation, a seller may choose *not ship* if she believes that this action will not be made known to future buyers. Even without information about a particular buyer’s feedback-provision history, she may assess the probability of receiving feedback by looking at her own feedback receipt history. In order to examine the effect of feedback receipt histories on seller decisions, we summarize these complex histories into simpler summary statistics.

In reviewing her feedback-receipt history, a seller may (and ought to) distinguish between feedback received after *ship* and *not ship*. Furthermore, she may distinguish between overall and more recent receipt. We define  $SellerNegRate$  ( $SellerPosRate$ ) as the proportion of a seller’s previous *not ship* (*ship*) decisions in the current subsession resulting in feedback. Summarizing recent feedback receipt, we define  $SellerSinceNeg$  ( $SellerSincePos$ ) as

---

<sup>22</sup>We assume that non-strategic reasons, such as altruism or reciprocity, are independent of the availability of buyer feedback histories.

<sup>23</sup>We also considered the alternative definition of a perfect reputation ( $perfectRep = 1$ ), i.e., those sellers with no negative feedbacks in the current subsession and at least one positive feedback. Thus, the set of sellers with a good reputation includes those with a perfect reputation. Considering the returns to a perfect reputation as opposed to a good reputation does not qualitatively alter our results.



the number of “missing” negative (positive) feedbacks since the seller last received negative (positive) feedback in the current subsession. Therefore, *SellerSinceNeg* equals 0 if the seller received feedback after the last time she did not ship, 1 if the seller did not receive feedback after the last time she chose *not ship*, but did receive feedback after her previous *not ship* decision, etc. *SellerSincePos* and *SellerPosRate* are missing if the seller has not previously chosen *ship* in the current subsession, as are *SellerSinceNeg* and *SellerNegRate* if she has not previously chosen *not ship* in the current subsession.

Sellers who have previously chosen both *ship* and *not ship* in the current subsession constitute a minority of shipping decisions in our experiment. We do not want to limit our analysis to only these observations. While these sellers have current information about feedback rates for both actions, such sellers might in fact be systematically different than other sellers. We would thus like some measure of a seller’s beliefs about the probability of feedback after each action, even for those sellers who have not chosen both actions in the current subsession. A seller in our experiment who has chosen *not ship* in a previous subsession might reasonably use this information in assessing the probability of feedback in the current subsession. To include these subjects in our analysis, we use the following measures. We define *SellerCumNegRate* (*SellerCumPosRate*) as the proportion of previous round 1–12 *not ship* (*ship*) decisions in any subsession resulting in feedback.<sup>24</sup> For example, in round 6 of the second subsession, we include rounds 1–12 of the first subsession, and the first 5 rounds of the second. We treat *SellerNegRate* as missing if the seller has not previously chosen *not ship* in rounds 1–12 of a subsession, and *SellerPosRate* as missing if she has not previously chosen *ship*. To summarize a seller’s most recent feedback receipt regardless of whether she has chosen both actions in the current subsession, we use *SellerSinceLastFB*, defined as the number of transactions since she last received feedback.

In Table 4, we present probit model estimation results with clustering at the individual level. In both models, the dependent variable is the indicator variable *Ship*. We control for the availability of buyer feedback histories by including the indicator variable  $D_{FBHist=1}$ . In these models, we lump together sellers who have chosen *not ship* at least once and possibly more than once in the current subsession with those who have always chosen *ship*. To control for this heterogeneity, we include *SellerPcntNoShip*, defined as ratio of *not ship* choices to sales in the current subsession.

In model (1), we include all observations in which the seller has previously chosen both *ship* and *not ship* before round 13 of a subsession. The estimated negative coefficient on the indicator variable for not selling an item in the most recent round in which the seller had a good reputation (*lastGoodNoSale*) is consistent with sellers using recent returns to a good reputation in assessing future returns to a maintaining or obtaining this reputation. The positive and significant coefficient on *SellerCumNegRate* suggests that past receipt of feedback after choosing *not ship* results in an increased probability of trustworthiness. In model (2), we limit our analysis to those sellers having chosen both actions in the current subsession. The coefficient on *lastGoodNoSale* remains negative and highly significant. The

---

<sup>24</sup>Defining these variables over all previous rounds does not qualitatively alter our results.

	(1)	(2)
	IsShip	IsShip
<b>Round</b>	-0.007	-0.005
	(0.004)	(0.011)
<b>D<sub>C=5</sub></b>	-0.035	-0.009
<i>=1 if C=5, 0 otherwise</i>	(0.028)	(0.060)
<b>D<sub>FBHist=1</sub></b>	-0.122	-0.140
<i>=1 if buyer history available, 0 otherwise</i>	(0.030)***	(0.052)***
<b>D<sub>subsess=3</sub></b>	0.077	0.088
<i>=1 if subsess=3, 0 otherwise</i>	(0.018)***	(0.037)**
<b>SellerPcntNoShip</b>	-0.203	-0.070
<i>Number No Ship/Number Sales</i>	(0.100)**	(0.234)
<b>lastGoodNoSale</b>	-0.363	-0.373
<i>isbuy<sub>τ</sub>, τ = max(t) s.t. goodRep<sub>t</sub> = 1</i>	(0.089)***	(0.130)***
<b>pcntGoodNoSale</b>	0.066	0.183
<i>1 - ∑<sub>t=1</sub><sup>round-1</sup> (isBuy<sub>t</sub> × goodRep<sub>t</sub>) / ∑<sub>t=1</sub><sup>round-1</sup> goodRep<sub>t</sub></i>	(0.110)	(0.298)
<b>SellerCumPosRate</b>	-0.091	
<i>Cumulative: Positive FBs/Number Ship</i>	(0.082)	
<b>SellerCumNegRate</b>	0.078	
<i>Cumulative: Negative FBs/Number No Ship</i>	(0.038)**	
<b>SellerSinceLastFB</b>	-0.007	
<i>Sales since seller's last feedback</i>	(0.006)	
<b>SellerPosRate</b>		-0.234
<i>Current Subsession: Positive FBs/Number Ship</i>		(0.090)***
<b>SellerNegRate</b>		0.101
<i>Current Subsession: Negative FBs/Number No Ship</i>		(0.089)
<b>SellerSinceLastFBPos</b>		-0.016
<i>Number Ships since last Positive FB</i>		(0.011)
<b>SellerSinceLastFBNeg</b>		-0.083
<i>Number No Ships since last Negative FB</i>		(0.106)
Observations	987	308
Log pseudo-likelihood	-391.89	-162.02

Notes: Coefficients are probability derivatives. Robust standard errors in parentheses are adjusted for clustering at the individual level. Significant at: \* 10-percent level; \*\* 5-percent level; \*\*\*1-percent level.

Table 4: Shipping in Rounds 2–12 in Subsessions 2 & 3: All Treatments: Probit Models with Clustering at Individual Level

coefficient on the rate of negative feedback receipt in the current subsession (*SellerNegRate*) has the expected positive sign, but is imprecisely estimated. We estimate a negative and significant coefficient on the rate of positive feedback receipt in the current subsession, which is consistent with sellers concern about negative feedback decreasing in positive feedbacks received.

In both models, we estimate a negative and significant coefficient on our indicator variable for availability of feedback-provision histories. Thus, even with reasonable controls for the value of a good reputation and a seller's own history of feedback receipt, there is less trustworthy behavior in our treatments with buyer feedback histories. We now focus on

these treatments to investigate how sellers used these histories.

To assess seller use feedback-provision histories, we summarize these complex histories into simpler summary statistics. As with our feedback-receipt summaries, we distinguish between overall and more recent feedback provision. However, whereas in “inspecting” her own history a seller can distinguish between feedback after *ship* and *not ship*, she cannot, in our experiment, discern whether feedback left by a buyer was positive or negative. Therefore, we define *BuyerFBRate* as the ratio of feedbacks left to purchases in the current subsession.<sup>25</sup> To summarize recent feedback provision, we define *BuyerSinceLastFB* as the number of feedback opportunities the buyer has missed since he last left feedback in the subsession. Therefore, *BuyerSinceLastFB* equals 0 if the buyer provided feedback after his last purchase, 1 if he did not provide feedback after his last purchase but did after the purchase immediately preceding his last purchase, etc.

In Table 5, we present probit model estimation results with clustering at the individual level. In all models, the dependent variable is the indicator variable *Ship*. Our data include all sessions in which a seller has access to the buyer’s feedback leaving history. In models (1), we include observation in which the seller has previously chosen both actions in any round 1–12. In model (2), we limit our observations to those in which a seller has previously chosen both actions in the *current* subsession. We control for a seller’s assessment of the returns to a good reputation with the indicator variable *lastGoodNoSale*, equal to 1 if the seller did not make a sale the last time she had a positive reputation. We also control for her assessment of receiving feedback based on her own history of feedback receipt as well as her buyer’s history of feedback provision. We control for heterogeneity specific to a seller in the particular subsession with *SellerPcntNoShip* (the ratio of *not ship* choices to sales in the current subsession).

Looking first at our larger subset of sellers having chosen both actions at least once in the past (model (1)), we estimate a negative and weakly significant coefficient on our indicator variable for high-cost sessions, and a positive and significant coefficient for the indicator variable for the last subsession in a session. The negative and significant coefficient on *lastGoodNoSale* suggests that assessments of the future returns to a good reputation continue to have an effect even after controlling for a seller’s feedback probability assessment.

In model (2), we restrict our attention to those sellers having chosen both actions in the current subsession, and are thus plausibly most confident in their assessment of feedback probabilities. The coefficient on the high-cost indicator variable is no longer significant, and our measure of the returns to a good reputation is now weakly significant as the coefficient is smaller in magnitude and less precisely estimated. Whereas the effect of the overall positive-feedback rate was weakly significant in model (1), the estimate of the effect of the current subsession’s positive-feedback rate in model (2) is larger and more precisely estimated. As the overall rate of positive feedback receipt includes both observable and no longer observable positive feedbacks, this is consistent with sellers less worried about

---

<sup>25</sup>Recall that sellers do not have access to a buyer’s feedback provision in earlier subsessions.

	(1)	(2)
	IsShip	IsShip
<b>Round</b>	-0.002	-0.006
	(0.008)	(0.016)
<b>D<sub>C=5</sub></b>	-0.100	0.032
<i>=1 if C=5, 0 otherwise</i>	(0.051)*	(0.072)
<b>D<sub>subsess=3</sub></b>	0.097	0.139
<i>=1 if subsess=3, 0 otherwise</i>	(0.025)***	(0.044)***
<b>SellerPcntNoShip</b>	-0.159	-0.006
<i>Number No Ship/Number Sales</i>	(0.179)	(0.269)
<b>SellerFBRate</b>	-0.015	
<i>Current Subsession: Feedbacks Received/Sales</i>	(0.073)	
<b>SellerCumPosRate</b>	-0.193	
<i>Cumulative: Positive FBs/Number Ship</i>	(0.107)*	
<b>SellerCumNegRate</b>	0.063	
<i>Cumulative: Negative FBs/Number No Ship</i>	(0.078)	
<b>SellerSinceLastFB</b>	-0.006	
<i>Sales since seller's last feedback</i>	(0.018)	
<b>BuyerFBRate</b>	0.012	-0.002
<i>Proportion of outcomes Buyer has revealed</i>	(0.075)	(0.143)
<b>BuyerSinceLastFB</b>	-0.035	-0.053
<i>Purchases since Buyer last revealed</i>	(0.011)***	(0.014)***
<b>lastGoodNoSale</b>	-0.294	-0.215
<i>isbuy<sub>τ</sub>, τ = max(t) s.t. goodRep<sub>t</sub> = 1</i>	(0.103)***	(0.119)*
<b>SellerPosRate</b>		-0.348
<i>Current Subsession: Positive FBs/Number Ship</i>		(0.105)***
<b>SellerNegRate</b>		0.045
<i>Current Subsession: Negative FBs/Number No Ship</i>		(0.113)
<b>SellerSinceLastFBPos</b>		-0.028
<i>Number Ships since last Positive FB</i>		(0.013)**
<b>SellerSinceLastFBNeg</b>		-0.161
<i>Number No Ships since last Negative FB</i>		(0.147)
Observations	505	204
Log pseudo-likelihood	-239.22	-105.10

Notes: Coefficients are probability derivatives. Robust standard errors in parentheses are adjusted for clustering at the individual level. Significant at: \* 10-percent level; \*\* 5-percent level; \*\*\*1-percent level.

Table 5: Round 2–12 Shipping in Subsessions 2 & 3 in Buyer Feedback History Treatments: Probit Models with Clustering at Individual Level

negative feedback the more current positive feedbacks.

In both models, while the coefficient on the overall rate of feedback provision (*BuyerFBRate*) is near zero and imprecisely estimated, we estimate a negative and significant coefficient on the number of transactions since the buyer last left feedback (*BuyerSinceLastFB*). Thus, even after including reasonable controls for both the value of a positive reputation and the seller's history of feedback receipt, we find that sellers base their shipping decision on buyers' feedback-provision histories. In particular, a seller is more likely to choose the untrustworthy action when faced with a buyer who has not *recently* provided feedback.

In sum, we find that a seller bases her decision both on the expected value of maintaining a reputation for trustworthiness and on the probability of feedback receipt. In assessing this probability, we find the a seller uses both her own history of feedback receipt as well as, when available, whether or not the buyer has not maintained a “reputation” for providing feedback.

### 3.3.2 Trust: The Purchase Decision

Buyers in our high-feedback-cost treatments were significantly less likely to purchase than their counterparts in either the low-feedback-cost treatments or the baseline sessions in which feedback was automatically provided. Lack of trust may be an indication of sellers not deserving trust. This seems particularly relevant for the high-cost sessions with buyer histories (*FBHistC5*) in which trustworthiness was significantly lower than in other treatments. However, this fails to explain the low levels of trust in the high-cost sessions without buyer histories (*NoHistC5*), which achieved levels of trustworthiness comparable to the low-cost and baseline sessions. We find that missing feedback contributes to a lack of trust and therefore helps explain the relatively low trust level in *NoHistC5*, the treatment with the lowest level of positive-feedback provision.

In deciding whether to trust, a buyer may use his own personal history of seller trustworthiness as well as available information about the current seller. We define two variables summarizing a buyer’s history with sellers. We define *BuyerPcntNoReceive* as the proportion of a buyer’s previous purchases in the current subsession in which the the seller chose *not ship*. To capture a buyer’s recent history, we set the indicator variable *BuyerLastExpNeg* equal to 1 if a seller did not ship the buyer’s most recent purchase, and equal to 0 otherwise.

We define analogous variables to summarize the seller’s public history. We define *SellerPcntSalesNeg* as the ratio of negative feedbacks received by a seller to the number of items sold in the current subsession. Due to missing feedback, *SellerPcntSalesNeg* does not give a complete summary of feedback received. We therefore define *SellerPcntSalesPos* as the ratio of positive feedbacks received by a seller to the number of items sold in the current subsession. To capture the seller’s recent actions, we set the indicator variable *lastFBNeg* equal to 1 if a is a seller’s most recent feedback is negative, and 0 otherwise.

In Table 6, we present probit model estimation results with clustering at the individual level. The dependent variable is the indicator variable *Buy* equal to 1 if the buyer chooses *buy*. We control for the availability of buyer feedback-provision histories with the indicator variable  $D_{FBHist=1}$  equal to 1 when available. We also include *BuyerPcntBuys*, defined as the buyer’s ratio of purchases to previous rounds in the current subsession, to capture any propensity to purchase specific to a buyer in a particular subsession. Finally, we investigate the possibility that untrustworthy behavior on the part of a seller believed to be trustworthy has a particularly chilling effect on that surprised buyer. We thus set the indicator variable *lastFBNegG* equal to 1 if the buyer’s last purchase was not shipped and that seller’s last

	(1)	(2)
	Buy	Buy
<b>Round</b>	-0.003	-0.007
	(0.002)	(0.003)**
<b>D<sub>subsess=3</sub></b>	0.019	0.004
<i>=1 if subsess=3, 0 otherwise</i>	(0.011)*	(0.013)
<b>D<sub>C=5</sub></b>	-0.010	-0.005
<i>=1 if C=5, 0 otherwise</i>	(0.021)	(0.027)
<b>D<sub>FBHist=1</sub></b>	0.029	
<i>=1 if buyer history available, 0 otherwise</i>	(0.020)	
<b>BuyerPcntNoReceive</b>	-0.139	-0.103
<i>Items Not Received/Items Bought</i>	(0.056)**	(0.091)
<b>BuyerPcntBuys</b>	0.236	0.161
<i>Purchases/(Round-1)</i>	(0.049)***	(0.065)**
<b>BuyerLastExpNeg</b>	-0.042	-0.046
<i>=1 if last purchase not received, 0 otherwise</i>	(0.027)	(0.033)
<b>SellerPcntSalesNeg</b>	-0.341	-0.451
<i>Negative FBs/Total Sales</i>	(0.099)***	(0.107)***
<b>SellerPcntSalesPos</b>	0.132	0.217
<i>Positive FBs/Total Sales</i>	(0.043)***	(0.052)***
<b>SellerLastFBNeg</b>	-0.220	-0.097
<i>=1 if last FB=Neg, 0 otherwise</i>	(0.046)***	(0.042)**
<b>BuyerLastExpNegG</b>	-0.080	-0.113
<i>BuyerLastExpNeg × (1-SellerLastFBNeg<sub>seller-1</sub>)</i>	(0.044)*	(0.067)*
<b>BuyerSinceLastFB</b>		0.000
<i>Purchases since Buyer last revealed</i>		(0.005)
Observations	2075	1035
Log pseudo-likelihood	-572.07	-278.46

Notes: Coefficients are probability derivatives. Robust standard errors in parentheses are adjusted for clustering at the individual level. Significant at: \* 10-percent level; \*\* 5-percent level; \*\*\*1-percent level.

Table 6: Buying in Rounds 2–12 of Subsessions 2 & 3: Probit Models with Clustering at Individual Level

recorded feedback at the time was positive.<sup>26</sup>

In model (1), we include as regressors variables capturing a particular buyer’s history with sellers in the current subsession, as well those capturing the information available about the current seller. The results strongly suggest a causal link from trustworthiness to trust. The negative and significant coefficients on *BuyerPcntNoReceive* (proportion of items bought but not shipped) and *SellerPcntSalesNeg* (proportion of sellers transactions resulting in negative feedback) provide evidence of this link and imply that buyers use both personal history as well as information about the particular seller. Furthermore, the negative and weakly significant coefficient on *lastFBNegG* is consistent with the hypothesis that the betrayal of trust from one believed to be trustworthy has a particularly chilling effect on

<sup>26</sup>Alternatively, we looked at the betrayal of trust by a seller without negative feedback at the time. The results are qualitatively similar.

trust.

Our results also shed light on the importance of missing feedback. The coefficient on *SellerPctSalesPos* (proportion of transactions resulting in positive feedback) is positive and significant. Thus, holding constant a seller’s negative feedbacks, transforming these missed feedbacks (a majority of which reflect positive outcomes) into positive feedbacks would increase the level of trust.

In Section 3.3.1 we presented evidence that sellers presented with feedback-provision histories discriminate against buyers who have not left feedback recently. A buyer fully cognizant of this would have incentives to either leave feedback more frequently or purchase less frequently. We would therefore expect (were feedback-provision histories randomly assigned) that a buyer who has not left feedback recently would be less likely to purchase. In model (2), we investigate whether this was the case. We include only observations from those sessions in which sellers had access to buyer feedback-provision histories. In addition to the regressors in model (4), we include *BuyerSinceLastFB*, the number of transactions since the buyer last left feedback. We find a very small and imprecisely estimated coefficient on *BuyerSinceLastFB*, suggesting that buyers did not to a large extent take their own feedback-provision histories into account when making their trust decision.

In sum, we find a strong link between trustworthiness and trust. While this link works through evidence of a particular seller’s trustworthiness, a buyer’s own experience informs his trust decision as well. In particular, a seller with a good reputation who acts in an untrustworthy manner reduces the displayed trust of that surprised buyer. Thus, the relatively low level of trustworthiness explains the relatively low levels of trust in the treatment with high feedback costs and availability of buyer feedback-provision histories (*FbHistC5*). We further find that missing feedback leads to lower levels of trust. Thus, the unwillingness of buyers to leave positive feedback in the *NoHistC5* treatment contributed to relatively low levels of trust despite a relatively high level of trustworthiness and rate of negative-feedback provision.

#### 3.4 Seller Discrimination and Feedback Provision

In light of buyers conditioning behavior on seller reputation, the seller strategy of discriminating against those who do not maintain a reputation for information sharing would not be wise if feedback-provision histories were not useful in predicting feedback receipt. Given the differing rates of positive and negative feedbacks, it is possible that feedback histories—determined in large part by feedback provision after trustworthy behavior by sellers—would not be useful in predicting feedback provision after *not ship*. In this section, while not ascertaining whether seller discrimination is optimal, we investigate whether sellers are in general correct in their belief that reticent sellers are less likely to give feedback after being cheated.

In Table 7, we present probit model estimation results with clustering at the individual

	(1)	(2)	(3)	(4)
	IsPosFB	IsNegFB	IsPosFB	IsNegFB
<b>Round</b>	0.002	0.02	0.003	0.024
	(0.005)	(0.011)*	(0.005)	(0.011)**
<b>D<sub>C=5</sub></b>	0.074	-0.148	0.061	-0.189
<i>=1 if C=5, 0 otherwise</i>	(0.061)	(0.097)	(0.061)	(0.097)*
<b>D<sub>subsess=3</sub></b>	-0.021	0.007	-0.016	0.048
<i>=1 if subsess=3, 0 otherwise</i>	(0.032)	(0.065)	(0.032)	(0.060)
<b>BuyerFBRate</b>	0.786	0.062	0.763	0.103
<i>Proportion of outcomes Buyer has revealed</i>	(0.109)***	(0.123)	(0.107)***	(0.120)
<b>BuyerSinceLastFB</b>	-0.024	-0.032	-0.024	-0.037
<i>Purchases since Buyer last revealed</i>	(0.015)	(0.018)*	(0.014)*	(0.018)**
<b>SellerCumNegRate</b>				0.041
<i>Cumulative: Negative FBs/Number No Ship</i>				(0.139)
<b>SellerSinceLastFB</b>			-0.006	0.016
<i>Sales since seller's last feedback</i>			(0.014)	(0.023)
<b>SellerCumPosRate</b>			0.068	
<i>Cumulative: Positive FBs/Number Ship</i>			(0.157)	
Observations	752	118	747	112
Log pseudo-likelihood	-213.80	-55.26	-212.58	-48.44

Notes: Coefficients are probability derivatives. Robust standard errors in parentheses are adjusted for clustering at the individual level. Significant at: \* 10-percent level; \*\* 5-percent level; \*\*\*1-percent level.

Table 7: Feedback Provision in Rounds 2–12 of Subsessions 2 & 3 in Buyer Feedback History Treatments: Probit Models with Clustering at Individual Level

level. The dependent variable is the indicator variable *isFB* equal to 1 if the buyer leaves feedback. In models (1) and (3), we include those observations in which the seller chose *ship*, whereas in models (2) and (4) we include observations where the seller chose *not ship*.

In models (1) and (2) we include as regressors those variables summarizing a buyer's feedback-provision history. Estimation results support these histories usefulness in predicting feedback. The positive and significant coefficient on *BuyerFBRate* in model (1) indicates that buyers who have given more overall feedback are in fact more likely to provide feedback should the seller choose *ship*. This makes sense, as high overall rates of feedback provision will be a result of high rates of positive feedback. The negative and weakly significant coefficient on *BuyerSinceLastFB* in model (2) provides evidence that buyers who have not recently given feedback are in fact less likely to provide negative feedback should the seller choose *not ship*. It therefore seems as though seller strategies discriminating against those who have not recently provided feedback are reasonable.

In order to address the possibility that these estimates are picking up a correlation between a buyer's feedback history and the overall rate of feedback provision, we include additional controls in models (3) and (4). As a proxy for overall levels of feedback provision, we include as regressors our variables summarizing a seller's recent receipt of feedback (*SellerSinceLastFB*) and overall rates of feedback receipt after *not ship* and *ship* (*SellerCumNegRate*



and *SellerCumPosRate*). We are in no way claiming a causal relationship between overall rates of feedback receipt and the likelihood of a particular buyer providing feedback. Rather, we are interested in assessing the robustness of a buyer’s feedback-provision history as a predictor of feedback provision when we include additional information that a seller might feasibly use in assessing the probability of feedback after each of her two actions. The estimates in models (3) and (4) provide evidence that the significance of buyer feedback-provision histories is robust to the inclusion of these other controls.

#### 4 DISCUSSION

In general, our results confirm that a reputation system can be effective in eliciting high levels of trust and trustworthiness even when sharing information about the trustworthiness of sellers is costly and voluntary. All of our treatments relying on the voluntary provision of this information performed about as well as our baseline in which this information was automatically provided.

However, our results suggest that in certain circumstances, additional information sharing may be beneficial. In particular, the level of trust in our treatment with a high cost of feedback provision and no availability of buyer feedback-provision histories lagged other treatments with similar levels of trustworthiness. We find that the low level of positive feedback in this treatment (and hence a high level of missing feedback) contributed to this lack of trust.

One mechanism which may increase feedback provision is the availability of buyer information-sharing histories. In our experiment, this did have the desired effect: it increased the rate of positive feedback when the cost of information sharing was relatively high and increased the rate of negative feedback when its cost was low. However, provision of this information allows sellers to target reticent buyers for untrustworthy behavior, a strategy unavailable when this information is unavailable. In fact, sellers in our experiment did just this. At least in this experiment, the end result of making available feedback-provision histories was no change in overall trust and trustworthiness, and a weakly significant decrease in trustworthiness in later subsessions. Our message to designers of e-commerce platforms seeking to entice information provision in this manner is to proceed with caution.

## REFERENCES

- Andreoni, James and John H. Miller**, “Rational Cooperation in the Finitely Repeated Prisoner’s Dilemma: Experimental Evidence,” *The Economic Journal*, May 1993, *103* (418), 570–585.
- Bakos, Yannis and Chrysanthos Dellarocas**, “Cooperation without Enforcement? A Comparative Analysis of Litigation and Online Reputation as Quality Assurance Mechanisms,” Working Paper 4295-03, MIT Sloan March 2003.
- Bolton, Gary E., Elena Katok, and Axel Ockenfels**, “How Effective Are Electronic Reputation Mechanisms? An Experimental Investigation,” *Management Science*, November 2004, *50* (11), 1587–1602.
- , — , and — , “Cooperation among Strangers with Limited Information about Reputation,” *Journal of Public Economics*, August 2005, *89* (8), 1457–1468.
- Dellarocas, Chris**, “The Digitization of Word of Mouth: Promise and Challenges of Online Feedback Mechanisms,” *Management Science*, October 2003, *49* (10), 1407–1424.
- Ely, Jeffrey and Juuso Välimäki**, “Bad Reputation,” *Quarterly Journal of Economics*, August 2003, *118* (3), 785–814.
- Fehr, Ernst and Simon Gächter**, “Cooperation and Punishment in Public Goods Experiments,” *American Economic Review*, September 2000, *90* (4), 980–994.
- Fischbacher, Urs**, “z-Tree: Zurich Toolbox for Ready-made Economic Experiments,” *Experimental Economics*, June 2007, *10* (2), 171–178.
- Gazzale, Robert S.**, “Giving Gossips Their Due: Information Provision in Games with Private Monitoring,” Economics Department Working Paper, Williams College August 2005.
- Johnson, Eric J., Colin F. Camerer, Sankar Sen, and Talia Rymon**, “Detecting Failures of Backward Induction: Monitoring Information Search in Sequential Bargaining,” *Journal of Economic Theory*, May 2002, *104* (1), 16–47.
- Kandori, Michihiro**, “Social Norms and Community Enforcement,” *Review of Economic Studies*, January 1992, *59* (1), 63–80.
- Kreps, David M. and Robert Wilson**, “Reputation and Imperfect Information,” *Journal of Economic Theory*, August 1982, *27* (2), 253–279.
- Li, Lingfang (Ivy)**, “Reputation, Trust, and Rebates: How Online Auction Markets Can Improve Their Feedback Mechanisms,” Research Paper No. 2008-05, Louisville College of Business October 2008.
- Resnick, Paul, Richard Zeckhauser, Eric Friedman, and Ko Kuwabara**, “Reputation Systems,” *Communications of the ACM*, December 2000, *43* (12), 45–48.

**Selten, Reinhard and Rolf Stoecker**, “End behavior in sequences of finite Prisoner’s Dilemma supergames A learning theory approach,” *Journal of Economic Behavior & Organization*, March 1986, 7 (1), 47–70.

**Siegel, Sidney and N. John Castellan Jr.**, *Nonparametric Statistics for the Behavioral Sciences*, 2nd ed., New York, NY: McGraw-Hill, 1988.

**Xiao, Erte and Daniel Houser**, “Emotion Expression in Human Punishment Behavior,” *Proceedings of the National Academy of Sciences of the United States of America*, May 17, 2005, 102 (20), 7398—7401.

## A EQUILIBRIUM PREDICTIONS

Our model follows Kreps and Wilson (1982). Our market (see Figure 1) is repeated  $N$  times, with markets indexed  $N, N-1, \dots, 2, 1$  (thus, Market  $N$  is the first market and Market 1 is the last market). We assume a sufficiently large number of buyers and sellers so that each buyer and seller are only matched once. We assume the existence of commitment sellers who always choose ship, and commitment buyers who might leave feedback even if not profit maximizing, with  $\kappa_\sigma$  and  $\kappa_\beta$  the proportions of commitment buyers and sellers. The remaining buyers and sellers are strategic, and thus only ship or leave feedback if it is in their best interest to do so. All buyers are strategic in their buying decisions.

We solve for sequential equilibria, with beliefs updated via Bayes' Rule. The belief that a seller is a commitment seller is  $\sigma$ , with  $\sigma_n^*$  the belief making a buyer in market  $n$  indifferent between buying and not buying. Likewise, the belief that a buyer is a commitment buyer is  $\beta$ , with  $\beta_n^*$  the belief making a seller in market  $n$  indifferent between shipping and not shipping. Let  $s_n$  be the probability a commitment seller ships given  $\sigma \geq \sigma_n^*$ ;  $f_n$  the probability a commitment buyer leaves feedback given  $\beta \geq \beta_n^*$ ; and  $b_n$  the probability a buyer buys in market  $n$ .

Under the assumption that a commitment seller with  $\sigma < \sigma_n^*$  does not ship, buying is at least weakly profit maximizing if

$$\underbrace{50\sigma_n + (1 - \sigma_n)(50s_n + 0(1 - s_n))}_{\text{buy}} \geq \underbrace{35}_{\text{not buy}},$$

$$\sigma_n > 1 - \frac{3}{10(1 - s_n)}, \text{ with}$$

$$\sigma_n^* = 1 - \frac{3}{10(1 - s_n)}. \tag{1}$$

If the fraction of commitment sellers is large, then by backwards induction buyers will always buy and strategic sellers will ship in all but the final market. We therefore make the following assumption about commitment sellers.

**Assumption 1.** *The fraction of commitment sellers is positive but less than  $\frac{7}{10}$  (i.e.,  $0 < \kappa_\sigma < \frac{7}{10}$ ).*

### A.1 No Feedback-Provision Histories

We first consider markets without feedback-provision histories. As each buyer-seller pair meets only once, strategic buyers will of course never leave feedback ( $f_n = 0$ ). We make the following general assumption about commitment buyers.

**Assumption 2.** *The probability with which a commitment buyer leaves feedback after a seller chooses not ship ( $f_\kappa^- > 0$ ) is at least as large as the probability with which a commitment buyer leaves feedback after a seller chooses ship ( $f_\kappa^+$ ).*

We believe that Assumption 2 is more realistic than the alternative assumption ( $f_{\kappa}^+ > f_{\kappa}^-$ ), but note that equilibrium predictions and restrictions are qualitatively similar under the alternative assumption. We simplify notation by defining the levels of commitment feedback:  $F_{\kappa}^+ = f_{\kappa}^+ \kappa_{\beta}$  and  $F_{\kappa}^- = f_{\kappa}^- \kappa_{\beta}$ . Note that when  $f_{\kappa}^+ = f_{\kappa}^- = \kappa_{\beta} = 1$ , we recover the equilibrium of the *Automatic* sessions.

As strategic buyers do not provide feedback, when  $f_{\kappa}^- = f_{\kappa}^+$ , missing feedback will be uninformative in the sense that  $\sigma_n = \sigma_{n-1}$  after missing feedback. When  $f_{\kappa}^- > f_{\kappa}^+$ , missing feedback is informative in the sense that  $\sigma_n \neq \sigma_{n-1}$  after missing feedback.

This informative of missing feedback adds a complication. Buyers provide incentives to ship by (probabilistically) buying after a “good outcome.” Letting  $\{\sigma_n^+, \sigma_n, \sigma^-\}$  be the beliefs after positive, missing and negative feedbacks, we have  $\sigma_n^+ > \sigma_n > \sigma_{n+1} > \sigma^- = 0$  for  $0 < s_n < 1$  and  $\sigma_{n+1} > 0$ . This leads to two possible equilibria. In the first, only positive feedback is “deemed” a good outcome, and the strategic seller chooses her shipping probability in Market  $n + 1$  so that  $\sigma_n^+ = \sigma_n^*$ . Buyers buy probabilistically after positive feedback, but do not buy after missing or negative feedback due to beliefs being too pessimistic. In the alternative equilibrium, the strategic seller chooses her shipping probability in Market  $n + 1$  so that  $\sigma_n = \sigma_n^*$ . Buyers always buy after positive feedback (as  $\sigma_n^+ > \sigma_n^*$ ), buy probabilistically after missing feedback, and do not buy after negative feedback.

We first look at the case where sellers choose  $s_n$  such that  $\sigma_{n-1}^+ = \sigma_{n-1}^*$ . With  $b_{n-1}^+$  the probability of purchase in market  $n - 1$  after a seller has received positive feedback, this  $s_n$  is a best response as long as

$$\underbrace{50 + F_{\kappa}^+ (70b_{n-1}^+ + (1 - b_{n-1}^+)35)}_{\text{ship}} + (1 - F_{\kappa}^+)35 \geq \underbrace{70 + 35}_{\text{not ship}},$$

$$b_{n-1}^+ \geq \frac{4}{7F_{\kappa}^+}, \text{ with} \tag{2}$$

$$b_{n-1}^{+*} = \frac{4}{7F_{\kappa}^+}$$

where  $b_1^{+*}$  is the probability of buying conditional on positive feedback that makes the seller indifferent. Equation (2) puts a lower bound ( $F_{\kappa}^+ \geq \frac{4}{7}$ ) on the probability of feedback after ship needed to support an equilibrium in which sellers choose  $s_n$  such that  $\sigma_{n-1}^+ = \sigma_{n-1}^*$ .

Assuming that buyers will only buy after positive feedback, the probability of commitment seller shipping in market  $n$  that leaves the  $n - 1$  buyer indifferent between buying and not is found by taking Bayes’ Rule and solving for  $s_n$ :

$$s_n = \frac{\sigma_n(1 - \sigma_{n-1}^*)}{(1 - \sigma_n)\sigma_{n-1}^*}. \tag{3}$$

To find  $\sigma_n^*$ , we plug equation (3) into equation (1) and solve recursively:

$$\sigma_n^* = (\sigma_1^*)^n = \left(\frac{7}{10}\right)^n.$$

We now look at the case where no news is good news, i.e., the market- $n$  seller has chosen  $s_n$  such that  $\sigma_{n-1} = \sigma_{n-1}^*$ . This  $s_n$  is a best response as long as

$$\underbrace{50 + 70F_\kappa^+ + (1 - F_\kappa^+)(70b_{n-1} + (1 - b_{n-1})35)}_{\text{ship}} \geq \underbrace{70 + F_\kappa^- 35 + (1 - F_\kappa^-)(70b_{n-1} + (1 - b_{n-1})35)}_{\text{not ship}},$$

$$b_{n-1} \geq \frac{4 - 7F_\kappa^+}{7(F_\kappa^- - F_\kappa^+)} = 1 - \frac{7F_\kappa^- - 4}{7(F_\kappa^- - F_\kappa^+)}, \text{ with} \quad (4)$$

$$b_{n-1}^* = 1 - \frac{7F_\kappa^- - 4}{7(F_\kappa^- - F_\kappa^+)}; \quad (5)$$

where  $b_{n-1}^*$  is the probability of making a sale in market  $n - 1$  that leaves the market  $n$  seller indifferent. Equation (4) puts restrictions on overall positive and negative feedback rates needed to support an equilibrium in which sellers choose  $s_n$  such that  $\sigma_{n-1} = \sigma_{n-1}^*$ . First, it must be the case that  $F_\kappa^+ < \frac{4}{7}$ . Intuitively, if the positive feedback rate is too high, the seller cannot sufficiently “punish” the buyer in the case of missing feedback. Likewise, it must be the case that  $F_\kappa^- > \frac{4}{7}$ .

In the case where buyers buy after missing feedback, we solve for the strategic seller’s probability of shipping that leaves the buyer in market  $n - 1$  indifferent:

$$\begin{aligned} \sigma_{n-1}^* &= \frac{\sigma_n (1 - F_\kappa^+)}{\sigma_n (1 - F_\kappa^+) + (1 - \sigma_n) (s_n (1 - F_\kappa^+) + (1 - s_n) (1 - F_\kappa^-))} \\ s_n &= \frac{\sigma_n (1 - F_\kappa^+) - \sigma_{n-1}^* (1 + (F_\kappa^- - F_\kappa^+) \sigma_n - F_\kappa^-)}{\sigma_{n-1}^* (F_\kappa^- - F_\kappa^+) (1 - \sigma_n)} \end{aligned} \quad (6)$$

Plugging equation (6) into equation (1) and solving recursively:

$$\sigma_n^* = \sigma_1^* \left( \frac{10 - 3F_\kappa^- - 7F_\kappa^+}{10(1 - F_\kappa^+)} \right)^{n-1} \quad (7)$$

**Proposition 1.** *In markets without feedback-provision histories, under Assumption 1, strategic sellers ship and buyers buy in the sequential equilibrium only if  $F_\kappa^- \geq \frac{4}{7}$ .*

**Proof:** We first establish the sequential equilibrium as a function of  $F_\kappa^+$  and  $F_\kappa^-$ . Consider the following strategy specification:

1. Strategic buyers never leave feedback.

2. Define  $\sigma_n^*$

$$\sigma_n^* = \begin{cases} 1 & \text{if } F_\kappa^- < \frac{4}{7} \\ \frac{7}{10} \left( \frac{10-3F_\kappa^- - 7F_\kappa^+}{10(1-F_\kappa^+)} \right)^{n-1} & \text{if } F_\kappa^- \geq \frac{4}{7} > F_\kappa^+ \\ \left( \frac{7}{10} \right)^n & \text{otherwise} \end{cases}$$

3. If  $\sigma_n > \sigma_n^*$ ,  $b_n = 1$ ; If  $\sigma_n < \sigma_n^*$ ,  $b_n = 0$

4. If  $\sigma_n = \sigma_n^*$ :

$$b_n = \begin{cases} 0 & \text{if } F_\kappa^- < \frac{4}{7} \\ \frac{4-7F_\kappa^+}{7(F_\kappa^- - F_\kappa^+)} & \text{if } F_\kappa^- \geq \frac{4}{7} > F_\kappa^+ \\ \frac{4}{7F_\kappa^+} & \text{otherwise} \end{cases}$$

5.  $s_1 = 0$ . For  $n > 1$ , if  $\sigma_n \geq \sigma_{n-1}^*$ ,  $s_n = 1$ . If  $0 < \sigma_n < \sigma_{n-1}^*$ ,

$$s_n = \begin{cases} 0 & \text{if } F_\kappa^- < \frac{4}{7} \\ \min \left\{ 1, \frac{\sigma_n(1-F_\kappa^+) - \sigma_{n-1}^*(1+(F_\kappa^- - F_\kappa^+)\sigma_n - F_\kappa^-)}{\sigma_{n-1}^*(F_\kappa^- - F_\kappa^+)(1-\sigma_n)} \right\} & \text{if } F_\kappa^- \geq \frac{4}{7} > F_\kappa^+ \\ \min \left\{ 1, \frac{\sigma_n(1-\sigma_{n-1}^*)}{(1-\sigma_n)\sigma_{n-1}^*} \right\} & \text{otherwise} \end{cases}$$

First, consider  $F_\kappa^- < \frac{4}{7}$ . By equations (2) and (4), there does not exist a  $b_1 \leq 1$  such that  $\pi(s_2 = 1) \geq \pi(s_2 = 0)$ . Therefore, strategic sellers will not ship in  $n = 2$ , and by backward induction in any  $n$ .

Now, consider  $F_\kappa^- \geq \frac{4}{7} > F_\kappa^+$ . Assume that a market-2 seller with  $0 < \sigma_2 < \sigma_1^*$  follows the proposed strategy. The seller will not ship in market 1 if she is strategic, and by equation (1), the seller's market-1 buyer is indifferent between buying and not if the seller does not receive market-2 feedback, and will buy if she received positive market-2 feedback. By equation (5), this seller is indifferent between shipping and not in market 2 if strategic. Note that shipping with a greater probability in market 2 is not an equilibrium ( $\rightarrow \sigma_1 < \sigma_1^* \rightarrow b_1 = 0 \rightarrow s_2 = 0$ , a contradiction), nor is shipping with a lower probability ( $\rightarrow \sigma_1 > \sigma_1^* \rightarrow b_1 = 1 \rightarrow s_2 = 1$ , a contradiction). Likewise, a larger  $b_1$  cannot be an equilibrium ( $\rightarrow s_2 = 1 \rightarrow \sigma_1 < \sigma_1^* \rightarrow b_1 = 0$ , a contradiction), nor can an increased probability of purchase after missing feedback be part of a sequential equilibrium, by a similar argument.

The logic of market-2 holds for markets  $n > 2$ .

Now, consider  $F_\kappa^+ \geq F_\kappa^- \geq \frac{4}{7}$ . Assume that a market-2 seller has received positive feedback in all  $n$  with  $s_n < 1$ , and has received no negative feedbacks. By the preceding argument, the proposed strategy is an equilibrium for  $n = 2$ , and by extension for  $n > 2$ . ■

## A.2 Feedback-Provision Histories

We start with the following assumption about the behavior of commitment buyers.

**Assumption 3.** *A commitment buyer leaves feedback after every purchase.*

We make this assumption to facilitate equilibrium analysis. Alternative assumptions about the behavior of commitment buyers would complicate the specification of strategies, but not alter the flavor of the equilibrium.

We show that feedback-provision histories enable equilibrium strategic-buyer shipping not available without these histories. Before specifying such a sequential equilibrium, we outline its structure. The strategic buyer does not leave feedback in Market 2 as Market-1 sellers do not condition actions on buyer feedback-provision histories. If the seller in Market 2 does not receive feedback, she does not sell in the Market 1. Thus, the Market-2 seller only ships (probabilistically) to the buyer who might be a commitment feedback provider. This provides the incentives for strategic buyer to masquerade as a commitment buyer.

Early in the game, the belief that the seller is commitment will be sufficient to induce buying and feedback. We then hit the point when the seller must randomize in order to make future sales. At this point, the strategic sellers choose  $s_n$  such that the *current* buyer is indifferent between leaving and not leaving *feedback*, while the strategic buyer randomizes so that the strategic seller is indifferent between shipping and not shipping. Note that the buyer with a positive reputation for leaving feedback will buy: if he was willing to pay  $C$ , it must be because he expects the buying in the next market to be worth at least  $35 + C$ . In Market 2, the strategic seller can no longer induce feedback provision from strategic buyers, and thus sets  $s_2$  to make the Market 1 buyer indifferent between buying and not buying, while the Market 1 buyer sets  $b_1$  to make the Market-2 strategic seller indifferent between selling and not selling.

While the intuition is fairly straightforward, the need to meet a partner with a “good” reputation complicates the analysis, especially if a seller (buyer) must assess the probability of meeting a good-reputation buyer (seller) in any future market as opposed to only in the next market. We thus solve only for the equilibrium in which a good buyer is one who has purchased and left feedback in every preceding market, while a good seller is one who has sold an item and received positive feedback in every preceding market.

To facilitate the analysis, we introduce new notation. First, we define  $S_n = \sigma_n + (1 - \sigma_n)s_n$  (the unconditional probability that a seller ships in market  $n$ ) and  $F_n = \beta_n + (1 - \beta_n)f_n$  (the unconditional probability a buyer provides feedback). We further define  $GS_{n-1} = GS_nGB_nS_nF_n$  (the probability of being matched with a “good” seller in market  $n - 1$ ) and  $GB_{n-1} = GS_nGB_nF_n$  (the probability of being matched with a “good” buyer in market  $n - 1$ ). Given  $\{GS_N, \dots, GS_n\}$ ,  $\{GB_N, \dots, GB_n\}$ , and the fractions of commitment sellers and buyers ( $\kappa_\sigma$  and  $\kappa_\beta$ ) we can recover the probabilities of strategic shipping and feedback.

In the final market, commitment sellers ship and strategic sellers choose *not ship*. Therefore



buying, regardless of feedback-provision history, is profit maximizing if:

$$\begin{aligned} \underbrace{50\sigma_1}_{\text{buy}} &\geq \underbrace{35}_{\text{not buy}} \\ GS_1 &\leq \frac{10\kappa_\sigma}{7}, \text{ with} \\ GS_1^* &= \frac{10\kappa_\sigma}{7} \end{aligned} \quad (8)$$

where equation (8) uses  $\sigma_n = \frac{\kappa_\sigma}{GS_n}$ . In order to make the market-1 buyer indifferent, the strategic market-2 seller chooses  $S_2^*$ :

$$\begin{aligned} GS_1^* &= GS_2GB_2F_2S_2 \\ &= GS_2GB_2\beta_2S_2 \\ &= GS_2S_2\kappa_\beta \\ S_2^* &= \frac{10\kappa_\sigma}{7GS_2\kappa_\beta} \end{aligned} \quad (9)$$

As market-1 sellers do not condition actions on buyer feedback-provision histories, market-2 strategic buyers do not provide feedback. Therefore, a market-2 seller will ship as long as

$$\underbrace{50 + \beta_2(70b_1 + (1 - b_1)35)}_{\text{ship}} + (1 - \beta_2)35 \geq \underbrace{70 + 35}_{\text{not ship}}$$

$$b_1 \geq \frac{4GB_2}{7\kappa_\beta}, \text{ with} \quad (10)$$

$$b_1^* = \frac{4GB_2}{7\kappa_\beta}, \quad (11)$$

where equation (10) uses  $\beta_n = \frac{\kappa_\beta}{GB_n}$ .

For markets  $n \geq 3$ , choosing ship is profit maximizing if

$$\underbrace{50 + 70F_nGB_{n-1} + (1 - F_nGB_{n-1})35}_{\text{ship}} \geq \underbrace{70 + 35}_{\text{not ship}} \quad (12)$$

$$GB_{n-1}^2 \geq \frac{4GB_nGS_n}{7}, \text{ with}$$

$$GB_{n-1}^* = \frac{2\sqrt{GB_nGS_n}}{\sqrt{7}}. \quad (13)$$

For markets  $n \geq 3$ , providing feedback is profit maximizing if

$$-C + GS_{n-1}S_{n-1}50 + (1 - GS_{n-1})35 \geq 35 \quad (14)$$

$$GS_{n-1} \geq \frac{GB_{n-2}C}{50GS_{n-2} - 35GB_{n-2}}, \text{ with}$$

$$GS_{n-1}^* = \frac{GB_{n-2}C}{50GS_{n-2} - 35GB_{n-2}}.$$

**Proposition 2.** *Under Assumptions and 1 and 3 in environments with feedback-provision histories, strategic sellers shipping and strategic-buyer feedback provision is a sequential equilibrium if  $\kappa_\sigma \geq \frac{7C\kappa_\beta}{500-245\kappa_\beta}$  and  $\kappa_\beta \geq \frac{8}{7\sqrt{7}}$ .*

**Proof:** Consider the following strategy specification:

1. Market-1 buyers buy with probability  $b_1 = \frac{4GB_2}{7\kappa_\beta}$ .
2. Market-1 strategic sellers do not ship ( $s_1 = 0$ ).
3. A buyer in market  $n \geq 2$  buys if he has provided feedback in every preceding market, and does not buy otherwise.
4. Market-2 strategic sellers ship with probability  $s_2 = \frac{(10-7\kappa_\beta)\kappa_\sigma}{7\kappa_\beta(GS_2-\kappa_\sigma)}$  if the buyer has provided feedback in every preceding market and she has received positive feedback in every preceding market, and does not ship otherwise.
5. Market-2 strategic buyers do not leave feedback.
6. A strategic buyer in market  $n \geq 3$  sets  $f_n$  such that  $GB_{n-1} = GB_{n-1}^*$ .
7. A seller in market  $n \geq 3$  ships with probability  $s_n$  such that  $GS_{n-1} = GS_{n-1}^*$  if she has received positive feedback in every preceding market and the buyer has left feedback in every preceding market. Otherwise, she does not ship.

We first establish the general structure of the sequential equilibrium. By equation (11), the market-2 strategic seller who has received positive feedback in every preceding market is indifferent between shipping and not shipping. By equation (9) the market-1 buyer is indifferent between buying and not buying. By the usual arguments, if  $b_1^* \leq 1$  and  $s_2^* \leq 1$ , no  $b_1 \neq b_1^*$  or  $s_2 \neq s_2^*$  can be part of a sequential equilibrium.

Assume that for all  $n \geq 4$ , all sellers have shipped and all buyers have bought and provided feedback. By definition,  $GS_3 = GB_3 = 1$ . By equation (13),  $GB_2^* = \frac{2}{\sqrt{7}}$ , and thus by equation (11)  $b_1^* = \frac{8}{7\sqrt{7}\kappa_\beta}$ . As  $b_1^* \leq 1$ ,  $\kappa_\beta \geq \frac{8}{7\sqrt{7}}$ . By equation (9) and inequality (14) with equality,  $GS_2^* = \frac{500\kappa_\sigma - 7C\kappa_\beta}{245\kappa_\beta}$ . Note that  $GS_2^* \geq \kappa_\sigma$  implies  $\kappa_\sigma \geq \frac{7c\kappa_\beta}{500-245\kappa_\beta}$ .

Note that it must be the case that  $GS_2 \leq GB_2$ . As  $F_2 = GB_2$ , it must therefore be the case that  $\kappa_\sigma \leq \frac{7\kappa_\beta}{10\sqrt{7}+C}$ , otherwise by equation (14) buyers leave feedback with certainty (and thus sellers ship with certainty). Furthermore, it must be the case that  $\kappa_\beta \leq GB_2 = \frac{2}{\sqrt{7}}$ , otherwise by equation (12) sellers always ship (and thus buyers always buy).

We now show that it cannot be the case that feedback and shipping is chosen probabilistically before market 3. Let  $m > 3$  be the first market in which buyers and sellers randomize.

First, it must be the case that  $GB_{m-1} \geq GS_{m-1}$ , or

$$\begin{aligned} GS_{m-1} &\leq GB_{m-1} \\ GS_{m-1} &\leq \frac{2}{\sqrt{7}}. \end{aligned} \tag{15}$$

Also,  $GB_{m-2} \leq GS_{m-1}GB_{m-1}$ , this means

$$\begin{aligned} GB_{m-2} &\leq GS_{m-1}GB_{m-1} \\ \frac{2}{\sqrt{7}}\sqrt{GS_{m-1}GB_{m-1}} &\leq GS_{m-1}GB_{m-1} \\ \frac{2^2}{7GB_{m-1}} &\leq GS_{m-1} \\ \frac{2}{\sqrt{7}} &\leq GS_{m-1} \end{aligned} \tag{16}$$

Inequalities (15) and (16) cannot generically hold at the same time. ■

## B EXPERIMENT INSTRUCTIONS

### Introduction

You are about to participate in a decision process in which one of numerous alternatives is selected in each of 45 rounds. This is part of a study intended to provide insight into certain features of decision processes.

During the session you will play a game that gives you an opportunity to earn cash. At the end of the session, you will be paid your earnings plus a \$3 show-up fee. If you make good decisions you may earn a considerable amount of money. Decisions and payments are confidential: No one will be told your actions or the amount of money you make.

### Procedures and Payoffs

You and the other participants in the room are the players in the game. Half of the players are randomly assigned to play the role of “buyers” whereas the other half play “sellers.” The experiment is divided into three sessions, with each session lasting 15 rounds. You will maintain the same role throughout the experiment.

In each round, each player is randomly matched with another player to trade a (fictional) commodity. First the buyer chooses to either *buy* or *not buy*. If the buyer chooses *not buy*, then the game ends and both players receive 35 points. If the buyer chooses to *buy*, then the game continues and the seller makes a decision to *ship* or *not ship*. *Ship* pays each player 50 points, while *not ship* pays the buyer 0 points and the seller 70 points.

After all sellers have made a decision, the buyers who have bought an item in the current round decide to either *leave feedback* or *not leave feedback*. Leaving feedback costs the buyer 5 points. If he chooses *leave feedback*, then all future buyers matched with this seller in this session **will** know whether or not this seller shipped the item in this round. If the buyer chooses *not leave feedback*, then future buyers matched with this seller in this session **will not** know whether or not the seller shipped the item in this round. If the seller did not get to move because the buyer did not buy from her, then the computer reveals the fact automatically.

Note that your earnings in each round depend only on the actions of you and the player with whom you are matched. Also, in each round, you are randomly matched with one of the players of the other role.

Figure 1 on the last page of the instruction illustrates the game procedure in a game tree. The exchange rate in the game is 1 dollar per 100 point.

### Information

Throughout the game, you will be able to view the information on your transaction history on the left side of the screen. The information about your current match will be displayed on the right side of the screen.

Each player can view his or her total earnings in the current session. The following set of information about a seller will be available to herself and her current match:

- the number of items sold (the number of times the buyer chose to buy from her) in the current session
- the number of positive and negative feedbacks received in the current session
- the number of rounds with no feedback in the current session
- For each round in the current session, a seller can see whether she sold the item, her shipping decision and whether feedback was left. A buyer can see a round-by-round history of the seller with whom he is currently matched: for each round in the current session, whether that seller sold an item, and, if the buyer in that round left feedback, whether the seller shipped the item.

The following set of information about a buyer will be available to himself and his current match:

- the number of times buying in the current session
- the number of times providing feedback in the current session
- For each round in the current session, a buyer can see whether he bought an item, whether the item was shipped to him, and whether he left feedback. A seller can see a round-by-round history of the buyer with whom she is currently matched: whether he bought an item and whether he left feedback for each round in the current session.

**The transaction information is reset in each session.** That is, after 15 rounds, a new session will start and **all** of the previous transaction information is no longer available. This means that your match cannot view any information about you from previous sessions. If

you wish to participate in this study, please read and sign the accompanying consent form. The consent form explains your rights as a subject as well as the rules of confidentiality that will be adhered to regarding your participation.

Round					Remaining time 19																																																																																																																																														
5 out of 15																																																																																																																																																			
Your current profile					YOU ARE A BUYER.					Seller's profile																																																																																																																																									
Number of Times Choosing to Buy 4										Number of Items Sold 2																																																																																																																																									
Number of Feedbacks Provided 2										Number of Positive Feedbacks 0																																																																																																																																									
Total Points Earned this Session 90										Number of Negative Feedbacks 1																																																																																																																																									
										Number of Sales Without Feedback 1																																																																																																																																									
<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th>Round</th> <th>Buy</th> <th>Ship</th> <th>Feedback</th> <th>Points</th> </tr> </thead> <tbody> <tr><td>1</td><td>Buy</td><td>Not Ship</td><td>No</td><td>0</td></tr> <tr><td>2</td><td>Buy</td><td>Not Ship</td><td>Yes</td><td>-5</td></tr> <tr><td>3</td><td>Buy</td><td>Ship</td><td>Yes</td><td>45</td></tr> <tr><td>4</td><td>Buy</td><td>Ship</td><td>No</td><td>50</td></tr> <tr><td>5</td><td></td><td></td><td></td><td></td></tr> <tr><td></td><td></td><td></td><td></td><td></td></tr> <tr><td></td><td></td><td></td><td></td><td></td></tr> <tr><td></td><td></td><td></td><td></td><td></td></tr> <tr><td></td><td></td><td></td><td></td><td></td></tr> <tr><td></td><td></td><td></td><td></td><td></td></tr> <tr><td></td><td></td><td></td><td></td><td></td></tr> <tr><td></td><td></td><td></td><td></td><td></td></tr> <tr><td></td><td></td><td></td><td></td><td></td></tr> <tr><td></td><td></td><td></td><td></td><td></td></tr> <tr><td></td><td></td><td></td><td></td><td></td></tr> <tr><td></td><td></td><td></td><td></td><td></td></tr> </tbody> </table>					Round	Buy	Ship	Feedback	Points	1	Buy	Not Ship	No	0	2	Buy	Not Ship	Yes	-5	3	Buy	Ship	Yes	45	4	Buy	Ship	No	50	5																																																												<p>The seller shipped your item.</p> <p>Please decide whether to provide feedback.</p> <input type="radio"/> Provide positive feedback <input type="radio"/> Do not provide feedback					<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th>Round</th> <th>Sold</th> <th>Feedback</th> </tr> </thead> <tbody> <tr><td>1</td><td>Sold</td><td>None</td></tr> <tr><td>2</td><td>Sold</td><td>Negative</td></tr> <tr><td>3</td><td>Not Sold</td><td>-</td></tr> <tr><td>4</td><td>Not Sold</td><td>-</td></tr> <tr><td>5</td><td></td><td></td></tr> <tr><td></td><td></td><td></td></tr> <tr><td></td><td></td><td></td></tr> <tr><td></td><td></td><td></td></tr> <tr><td></td><td></td><td></td></tr> <tr><td></td><td></td><td></td></tr> <tr><td></td><td></td><td></td></tr> <tr><td></td><td></td><td></td></tr> <tr><td></td><td></td><td></td></tr> <tr><td></td><td></td><td></td></tr> <tr><td></td><td></td><td></td></tr> </tbody> </table>					Round	Sold	Feedback	1	Sold	None	2	Sold	Negative	3	Not Sold	-	4	Not Sold	-	5																																
Round	Buy	Ship	Feedback	Points																																																																																																																																															
1	Buy	Not Ship	No	0																																																																																																																																															
2	Buy	Not Ship	Yes	-5																																																																																																																																															
3	Buy	Ship	Yes	45																																																																																																																																															
4	Buy	Ship	No	50																																																																																																																																															
5																																																																																																																																																			
Round	Sold	Feedback																																																																																																																																																	
1	Sold	None																																																																																																																																																	
2	Sold	Negative																																																																																																																																																	
3	Not Sold	-																																																																																																																																																	
4	Not Sold	-																																																																																																																																																	
5																																																																																																																																																			
<div style="border: 1px solid black; display: inline-block; padding: 5px 10px; background-color: #f00; color: white; margin: 10px 0;">Submit</div>																																																																																																																																																			

Figure 4: A typical buyer's screen in a high-cost treatment.

C COMPUTER SCREEN