The B.E. Journal of Economic Analysis & Policy

Contributions

Volume 10, *Issue* 1 | 2010 | *Article* 92

I Scratched Yours: The Prevalence of Reciprocation in Feedback Provision on eBay

Lian Jian* Jeffrey K. MacKie-Mason† Paul Resnick‡

*University of Southern California, ljian@umich.edu
†University of Michigan, jmm@umich.edu
‡University of Michigan, presnick@umich.edu

Recommended Citation
Lian Jian, Jeffrey K. MacKie-Mason, and Paul Resnick (2010) "I Scratched Yours: The Prevalence of Reciprocation in Feedback Provision on eBay," *The B.E. Journal of Economic Analysis & Policy*: Vol. 10: Iss. 1 (Contributions), Article 92.
Available at: http://www.bepress.com/bejeap/vol10/iss1/art92

I Scratched Yours: The Prevalence of Reciprocation in Feedback Provision on eBay*

Lian Jian, Jeffrey K. MacKie-Mason, and Paul Resnick

Abstract

Many online systems for bilateral transactions elicit performance feedback from both transacting partners. Such bilateral feedback giving introduces strategic considerations. We focus on *reciprocity* in the *giving* of feedback: how prevalent a strategy of giving feedback is only if feedback is first received from one's trading partner. The overall level of feedback activity clearly depends on the prevalence of the reciprocation strategy: in a market with many reciprocators and few unconditional feedback providers, the equilibrium quantity of feedback can be quite low. We estimate the prevalence of such reciprocation in one market, eBay. Reciprocation cannot be directly distinguished from late feedback that was not conditioned on the partner having provided feedback. We develop a model that distinguishes the two by exploiting information about the timing of feedback provision when the partner does not provide feedback. We find that buyers and sellers on eBay used the "reciprocate only" strategy about 20-23% of the time. We also measure the extent to which the prevalence of these strategies changes with the experience levels of the two parties and with the item price.

KEYWORDS: feedback provision, reputation systems, e-commerce

*We are grateful for Tapan Khopkar's earlier participation in developing this paper. We also appreciate feedback at various stages of this study from Gary Solon, Chris Dellarocas, Nigel Melville, Sijian Wang, and participants of the ICD lab group at the University of Michigan (Rick Wash, Anna Osepayshvili, John Lin, Joshua Cherry, and Ben Stearns). This material is based upon work supported by the National Science Foundation under Grant No. IIS 0308006. Lian Jian was also supported by the School of Information at the University of Michigan and by the Rackham Predoctoral Fellowship while completing the study. We thank eBay for providing the data used in this analysis.

1 Introduction

Reputation systems enable trade among strangers by informing people about their trading partner's past performance, which also creates incentives for good behavior (Resnick, Zeckhauser, Friedman, and Kuwabara, 2000, Dellarocas, 2003). Many online marketplaces offer reputation systems based on subjectively provided feedback. For example, buyers rate sellers on Amazon's marketplace, and those ratings are visible to future buyers.

In some cases, feedback provision is two-sided. For example, at couchsurfing.net, where travelers find free places to stay while traveling, both hosts and travelers can rate each other. eBay buyers and sellers can rate each other. Many other two-sided markets are also candidates for two-sided reputation systems, such as dating sites, housing matches, and ride-sharing services.

Subjectively reported two-sided feedback introduces strategic considerations: whether to provide feedback, and the content of that feedback, may be influenced by the partner's actual or expected feedback-giving behavior. For example, anecdotes suggest that some eBay users employ a feedback-giving strategy we call "reciprocation": they only give feedback after receiving feedback from their trading partners.[1] Anecdotal evidence also suggests that buyers and sellers withhold negative feedback in order to avoid receiving retaliatory negative feedback. eBay conducted experiments on alternative feedback designs and, in 2008, removed the option for sellers to provide negative feedback to buyers, though they can still provide positive feedback and buyers can still provide either kind (Bolton, Greiner, and Ockenfels, 2009).

While two-sided subjective feedback may inhibit provision of negative feedback, it may help solve an under-provision problem for positive feedback. Feedback information is a public good: one person's consumption of published feedback does not diminish another's use of it. Theory predicts that in general public goods will be under-provided (Samuelson, 1954). The free-rider problem seems especially pernicious because feedback only benefits other users, not its provider, so that self-interested users then appear to have little or no incentive to provide feedback. Nevertheless, more than half of the traders on eBay provide feedback (Resnick and Zeckhauser, 2002). Why?

[1] For example, on Yahoo!Answers a user posted the question "Why do eBay sellers not give feedback as soon as you pay?" and received the following best answer: "Many sellers wait until they receive feedback from the buyer before they leave feedback" (http://answers.yahoo.com/question/index?qid=20060816131534AAecgGz, retrieved on Sept 14, 2010). Similar conversations also occurred on eBay's forum (e.g., http://reviews.ebay.com/Who-Should-Leave-Feedback-First-The-Buyer-or-Seller_W0QQugidZ10000000000766286, retrieved on Sept 14, 2010).

There are many possible motivations, likely experienced to a greater or lesser degree by different people. For example, some may freely contribute feedback because they are altruists, willing to incur a small cost to contribute to the community (Fehr and Gächter, 2002). Some may exhibit "reciprocal altruism" (Andreoni and Miller, 1993, Gächter and Falk, 2002), a tendency to give people "what they deserve", in this case a positive feedback in return for a good transaction and a negative feedback in return for a bad one. Some may provide feedback to avoid the hassle of partners asking for or demanding it. Some may fear that if word gets around that they don't provide feedback, partners will take advantage of them (Gazzale, 2004, Chapter 1).

Another reason to provide feedback is that it may spark the desirable event of the partner providing feedback. If many of one's transaction partners employ a feedback reciprocation strategy, then providing feedback first can be a way to build one's own feedback profile faster. If everyone followed a strategy of reciprocating feedback, and no one chose to give unconditionally, no one would ever be the first to provide feedback and none would be provided. On the other hand, if no one followed a reciprocation strategy, one of the incentives for providing feedback would be removed. Estimating the prevalence of feedback reciprocation provides a window into the complex ecology of feedback provision.

Of course, that begs the question of why a self-interested party would employ a strategy of giving feedback after receiving it. Clearly, having one's partner expect such feedback reciprocation can induce the partner to provide positive feedback (in order to get it in return) and to remain silent rather than providing negative feedback after a bad transaction (in order to avoid getting it in return). In a one-shot game, however, in the absence of a binding feedback reciprocation contract, actually delivering the reciprocal feedback, at some cost of effort, would not be rational. With repeated interactions, some form of direct retaliation may be sufficient for a reciprocation equilibrium. Or, even without repeated direct interaction, a generalized reciprocity equilibrium may emerge (Jian and MacKie-Mason, 2008). As with any feedback provision, some reciprocators may also want to follow through on providing feedback for the non-rational reason of giving a gift to, or taking vengeance on their partners, in this case rewarding or punishing the partners' feedback rather than the partners' action in the underlying sales transaction.

We do not propose a specific theoretical model for why feedback reciprocation occurs: we would not be able to estimate a structural model with the data available to us in any case. Rather, using a large dataset of eBay transactions we test for the prevalence of reciprocation, and the prevalence of two alternative strategies, unconditional provision and non-provision. We find that buyers and sellers on eBay used the "reciprocate" strategy about 20-30% of the time. We also measure the extent to which the prevalence of these strategies changes with the experience levels of the two parties, and with the item price. For example, in bilateral transactions, the relative experience levels may matter as inexperienced traders learn the strategy equilibrium, while more experienced traders may be trying to teach their partners.[2] Dellarocas and Wood (2008) have found that the level of traders' satisfaction varies with item price; we explore whether it affects the prevalence of reciprocal or non-reciprocal feedback giving strategies.

We also make methodological contributions to the estimation of feedback provision strategies in two-sided reputation systems. When both parties provided feedback, it is not clear whether they did so independently or whether the second did so in response to the first. When neither party provided feedback, it is not clear whether their decisions were unconditional or whether one or both would have provided feedback had the other done so. This is a problem in estimating choice when the underlying decision variables are latent (unobservable). We develop a latent variable estimation procedure that takes advantage of the observable timing of feedback provision to identify and estimate reciprocal feedback-giving strategies. We are not the first to develop econometric models to identify reciprocal feedback-giving strategies. Previously, Dellarocas and Wood (2008) developed a different model to study both the biases in feedback ratings and reciprocal feedback-giving behavior. Our models differ in various ways, as detailed in section 2.

2 Related Work

Previous work on feedback provision has estimated the impact of prior negative feedback in a seller's history on the buyers' willingness to provide negative feedback. Resnick and Zeckhauser hypothesized that buyers may "stone" sellers who have received negative feedback, becoming harsher in their assessments of later transactions (Resnick and Zeckhauser, 2002). Empirically, the probability of receiving a negative feedback increases immediately following the receipt of one (Cabral and Hortaçsu, 2010). There are several possible explanations besides stoning, including slipping (the first and subsequent negatives were both the result of the same decline in seller quality) and slacking (the seller provides lower quality because of receiving the first negative). By modeling the different but overlapping time windows in which these different explanations would operate, Khopkar, Li, and Resnick (2005) showed that some of the effect is indeed due to stoning. This line of work does not address the role of feedback reciprocation between partners to a given transaction.

[2]In another context, Wikipedia has a "welcoming committee" in charge of greeting new members, introducing the community's policies, guidelines, and social norms to them (Wikipedia, 2010).

A few studies have provided empirical evidence for the *existence* of strategic feedback reciprocation by exploring the correlations between buyers' and sellers' feedback timing (Bolton et al., 2009, Dellarocas and Wood, 2008, Resnick and Zeckhauser, 2002), but none of these studies offers an estimate of the *prevalence* of strategic feedback reciprocation. Bolton et al. (2009) also conducted human-subject experiments to compare the effects of various feedback provision mechanisms on the efficiency of the electronic market.

Dellarocas and Wood (2008), hereafter DW, is closest in spirit to our analysis. DW have two main results: a *feedback bias* result — traders report different transaction outcomes (positive, neutral, and negative) with different probabilities, leading to biases in the aggregated probabilities of various outcomes — and a *feedback reciprocation* result — that feedback received increases the probability of feedback giving. Like DW, we identified the existence of feedback reciprocation among eBay traders. We go further and calculate the magnitude of reciprocation, and we also measure the extent to which some observable factors influence traders' choices to strategically reciprocate feedback.

One key difference between our models is that we make different assumptions on how the timing of the first feedback affects the likelihood that the receiver reciprocates (i.e., changes her action, providing a feedback that she otherwise would not have.) We assume the proportion of reciprocators in the population is unaffected by the timing of the partner's feedback, whereas in DW's model, the number of reciprocators declines the longer the partner waits to provide the first feedback. In both models, there are fewer people left to go second the longer it takes for the first feedback to arrive, because more of those who were going to give feedback whether they received it or not will already have given feedback. In our model, the number of people who *change* their action to one of feedback giving is unaffected by when the partner provides feedback.

Another major difference between our models is that our method does not require a parametric assumption on how the receipt of the first feedback affects the timing distribution of the second feedback. DW estimate whether there is reciprocation by estimating whether receiving feedback increases the partner's subsequent hazard rate for a particular time-to-feedback distribution (Dellarocas and Wood, 2008, Section 4.1.2).[3] By contrast, we make parametric assumptions only about the timing distribution of feedbacks when the partner does *not* give feedback. From that, we estimate the total probability mass of feedback expected to be given after the partner's feedback, without making a parametric assumption about its timing distribution.

[3] The hazard rate $h(t)$ of a failure time distribution $F(t)$ is defined as the rate of an event occurring given that it didn't happen in the past. That is, $h(t) = F'(t)/(1-F(t))$.

One benefit of our approach is that we are then able to estimate how much various observable factors affect the *prevalence* of a feedback reciprocation strategy. In particular, we estimate how much the traders' prior feedback profiles, and the item price, affect the traders' feedback giving strategy choices. In principle, DW's model could be extended for similar analyses, but it would be less straightforward to interpret. Instead of estimating the impact of a one dollar increase in item price on the percentage of buyers pursuing a reciprocation strategy, the natural extension of DW's model would yield an estimated impact on the hazard rate. For a particular time of first feedback, this could be rolled up into a total expected number of extra (reciprocating) second feedbacks, but not an estimated impact that is independent of the time of the first feedback.

3 Model Description

3.1 Feedback outcomes and strategies

We do not observe eBay traders' feedback provision strategies (that is, their internal mental plans), only their observable feedback-giving actions. All transactions on eBay result in one of the following five outcomes in terms of feedback provision: No Feedback, only the seller gives feedback (Seller Only), only the buyer gives feedback (Buyer Only), the seller gives feedback first and the buyer next (Seller First), or the buyer gives feedback first and the seller next (Buyer First). Any distribution of unobserved strategies would generate a distribution of these five observable outcomes. We develop an estimation strategy that allows us to econometrically identify the latent (unobservable) distribution of strategy choices, given the observable outcomes.

We posit that each trader adopts, on each transaction, one of three strategies: abstain from giving feedback (N), give unconditionally (Y), or reciprocate (R). Y means that the player gives feedback on the transaction regardless of whether the partner does. R means that the player gives feedback on that transaction only if, and only after, the partner does.

In our stochastic model, we express buyers' and sellers' strategies in probability terms. Let R_g denote the probability that a trader with role $R \in \{B,S\}$ plays strategy $g \in \{y,r,n\}$. For example B_y is the probability that a Buyer plays strategy Y of giving feedback unconditionally. Thus, we can think of (B_y, B_r, B_n) as a mixed strategy that a buyer will follow on a particular transaction. The mix may depend on the item price, the number of prior feedbacks each partner has received, and many unobserved and unmodeled characteristics of the transaction and the buyer.

Outcome	**Buyer Strategy**	**Seller Strategy**
No Feedback	N or R	N or R
Seller Only	N	Y
Buyer Only	Y	N
Seller First	Y or R	Y
Buyer First	Y	Y or R

Table 1: Mapping Strategies To Feedback Outcomes

Intuitively from Table 1 it is clearly not possible to identify statistically the prevalence of the three strategies merely from observing which of the five feedback outcomes is realized after a transaction. Immediately, from the first row, if no feedback was provided by either the seller or the buyer, we know that both chose to either abstain or reciprocate, but we cannot tell which. However, in some transactions we can identify the strategy from the outcome. For example, if only the seller (buyer) gave feedback (second and third rows), it must have been that the seller (buyer) chose the give unconditionally strategy, and the buyer (seller) choose to abstain. Last, if both the seller and the buyer gave feedback, with one first and the other later (fourth and fifth rows), we can infer the strategy of only one of them: the first giver must have chosen to give unconditionally, but the second could have chosen either an unconditional giving strategy Y (but happened to act slower than the first provider), or a reciprocation strategy R.

3.2 Feedback timing

Although we cannot uniquely identify strategies from outcomes alone, we have more information available to us: we have the time at which feedback was provided. In our dataset, for each feedback we observe the time at which it was given, denoted by t_s if given by the seller and t_b if by the buyer, expressed as offsets or time elapsed from the close of bidding on the auction of the item. This information is sufficient to enable us to identify the strategies. For example, notice that the probability of observing an outcome in which both provide feedback, but the buyer gives first, depends on the timing:

$$Pr(\text{Buyer First}) = B_y \cdot S_y \cdot Pr(t_s > t_b | B_y, S_y) + B_y S_r, \tag{1}$$

where t_s and t_b are the time at which the seller and the buyer give feedback respectively. Thus, if we could estimate $Pr(t_s > t_b | B_y, S_y)$, it would help in identifying the quantities B_y, S_y, and S_r.

Conditional on the trader's role, $r \in \{b, s\}$, and her strategy choice, $g \in \{y, r\}$, we assume her time of feedback, t_r, follows a distribution described by the probability density function $f_{rg}(t_r)$. For example, for a seller who plays the unconditional strategy Y, the probability that she gives feedback at time t_s is $f_{sy}(t_s)$. We assume the timing distributions of both buyers' and sellers' feedback are lognormal, and write the feedback timing distribution for sellers playing the Y strategy as $f_{sy}(t_s) = LNORM(t_s)$, and similarly, for buyers playing the Y strategy as $f_{by}(t_b) = LNORM(t_b)$.[4] We can obtain an unbiased estimate for f_{sy} (f_{by}) using only observations with the Seller (Buyer) Only outcome.

In order to provide intuitions on how we use feedback timing to separately identify unconditional feedback givers and reciprocators, we illustrate our model of feedback timing in Figure 1. Imagine a dataset containing lots of transactions with the buyer giving feedback first, all at time t_b. Some of the sellers who subsequently gave feedback were following an unconditional strategy Y, while other may have been following a reciprocation strategy R. For those following strategy Y, absent buyer feedback the probability density of seller feedback at time t_s is $f_{sy}(t_s)$. Even with the buyer feedback, among sellers following strategy Y, behavior before time t_b should follow the same distribution, and thus a fraction $\alpha = F_{sy}(t_b)$ should have given feedback before t_b, and only a fraction $(1-\alpha)$ are left to give feedback after t_b. In other words, $Pr(t_s > t_b | B_y, S_y) = 1 - \alpha = 1 - F_{sy}(t_b)$. If there were no sellers following reciprocation strategy R, and we had estimates of B_y and S_y, the overall probability of the buyer giving feedback first, at time t_b, would be $B_y f_{by}(t_b) S_y (1 - \alpha)$. If the dataset shows that the probability of a seller giving feedback after t_b is greater than $S_y(1-\alpha)$, the extra must have come from the reciprocators, shown as the shaded area marked with S_r.

Figure 1 shows a specific case with Buyer First outcome, in which the buyer gives feedback at time t_b. In our dataset, t_b, as well as t_s, can vary across the whole time axis, e.g., $t_b^{'}$ in Figure 1, and we do not have the luxury of many transactions for each particular value of t_b. Using the same logic illustrated in Figure 1, however, any values of *LNORM* parameters defining f_{sy} and f_{by} will determine a likelihood of each of the transaction observations, with their actual t_s and t_b when feedbacks are provided. Thus, maximum likelihood estimation can be used to select *LNORM* parameters that best fit the observed data.

[4] Section 4.2 and appendix A explore alternative functional forms besides the lognormal distribution.

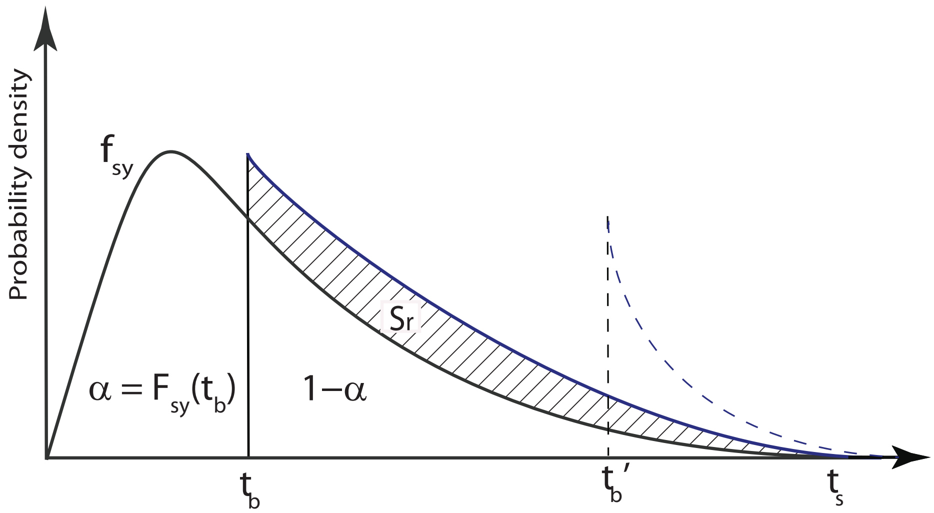

Figure 1: The feedback timing model with the Buyer First outcome.

3.3 The likelihood function

With the definition of feedback outcomes, strategy space, and timing distribution function, we construct a multinomial maximum likelihood model with simultaneous equations. Let θ denote the vector of parameters to be estimated, which will be explained in section 3.4. Equation (2) is the overall likelihood function of θ given all the observable response variables Z in our dataset. For each transaction i, Z_i consists of the feedback provision outcome m_i, $m_i \in M$ where $M = \{No\ Feedback, Seller\ Only, Buyer\ Only, Seller\ First, Buyer\ First\}$, and the times of those feedbacks, if they occur. Our likelihood function is as follows,

$$L(\theta;Z) = \prod_{i=1}^{N} l(\theta;m_i,t_b,t_s) \tag{2}$$

The outcome of No Feedback is observed whenever neither the seller nor the buyer played the Y strategy (giving unconditionally). Thus, the likelihood of θ for No Feedback observations is,

$$l(\theta;No\ Feedback) = (1-S_y)(1-B_y), \tag{3}$$

When the outcome of Seller Only is observed, the likelihood is

$$l(\theta;Seller\ Only,t_s) = B_n S_y f_{sy}(t_s). \tag{4}$$

Similarly, when the outcome is Buyer Only, the likelihood is

$$l(\theta; Buyer\ Only, t_b) = S_n B_y f_{by}(t_b). \tag{5}$$

If the outcome of Buyer First is observed, the likelihood is as follows,

$$l(\theta; Buyer\ First, t_b, t_s) = \tag{6}$$
$$\underbrace{B_y f_{by}(t_b)}_{A} \left(\underbrace{S_y}_{B} \left(\underbrace{Pr(t_s > t_b | B_y, S_y)}_{C} \right) + \underbrace{S_r}_{D} \right),$$

where term A contains the probability that the buyer was playing the strategy Y (give feedback unconditionally), and that he chooses this particular time, t_b, to give feedback; Terms B and D contain the probabilities that the seller might be playing the Y strategy and the R strategy respectively. Term C specifies the probability that if the seller is playing the Y strategy, her time of feedback happens to be later than the buyer's. Using F_{sy} to denote the corresponding cumulative distribution function of f_{sy}, we can write the Term C as:

$$Pr(t_s > t_b | B_y, S_y) = 1 - F_{sy}(t_b) \tag{7}$$

Similarly, if the outcome of Seller First is observed, the likelihood is as follows,

$$l(\theta; Seller\ First, t_b, t_s) = S_y f_{sy}(t_s) \left(B_y Pr(t_b > t_s | B_y, S_y) + B_r \right), \tag{8}$$

where $Pr(t_b > t_s | B_y, S_y) = 1 - F_{by}(t_s)$.

3.4 Functional form assumptions

To estimate parametrically the probabilities of trader i playing each of the three strategies, we make the assumption that her probabilities of choosing any one of the three strategies are governed by multinomial logistic distributions:

$$\begin{aligned} B_{yi} &= \frac{e^{\beta_y X_i}}{1 + e^{\beta_y X_i} + e^{\beta_r X_i}}, \\ B_{ri} &= \frac{e^{\beta_r X_i}}{1 + e^{\beta_y X_i} + e^{\beta_r X_i}}, \end{aligned}$$

$$
\begin{aligned}
S_{yi} &= \frac{e^{\theta_y X_i}}{1+e^{\theta_y X_i}+e^{\theta_r X_i}}, \\
S_{ri} &= \frac{e^{\theta_r X_i}}{1+e^{\theta_y X_i}+e^{\theta_r X_i}},
\end{aligned}
$$

where X_i is the vector of independent variables that we will later use in the regression. B_{ni} and S_{ni} can be derived from the above expressions.

3.5 Model validation

To validate the mathematical model and our STATA code, we conducted Monte Carlo simulations. We generated datasets according to our functional form assumptions and a set of arbitrarily chosen parameters. See appendix C for the true values of all the parameters used in the simulation, and the estimates for these parameters found using our model. Summarizing our simulation results, the distribution of simulation errors is quite close to the predicted distribution. In the last column of Table 14 we report the errors in units of standard deviations, and in Figure 6 we plot the cumulative distribution of these errors against the asymptotic normal distribution the Monte Carlo should generate. The match is quite good, and there are no outliers. We conclude that our model is well identified and correctly programmed.

4 Data Set

We derived our sample from three master datasets provided by eBay:

1. Items Dataset: contains transactional data for all the items listed for sale on eBay from February 1st 1999 to June 30th 1999.
2. Feedback Dataset: contains all feedback data up to May 31st 1999.
3. Users Dataset: contains the id and registration dates for all the users who registered before June 30th 1999.

Some buyer-seller pairs conduct multiple transactions, and feedback giving patterns may be quite different on subsequent transactions than on initial transactions, especially since eBay did not count multiple feedbacks from the same partner in a trader's score, thus potentially changing the incentive to provide multiple feedbacks to the same partner. Moreover, at the time of our dataset, eBay did not require that all feedback be tied to a specific transaction. Thus, we chose the buyer-seller partnership's first transaction as the unit of analysis.

We extracted all the items listed for auction during the first week of March 1999 and eventually purchased, involving buyer-seller partnerships that had not conducted a prior transaction and had no prior feedback. This initial sample contains 959,657 items.[5]

The auction of an item ends with a winning bidder, whose bid is higher than other bidders' bids and the reservation price set by the seller. From our perspective, this marks the start of a transaction.

At the time of these transactions, eBay opened the feedback channels for both buyer and seller to rate each other as soon as the transactions started. Each feedback contained two parts: an indicator (+1 for positive, -1 for negative, or 0 for neutral) and an optional text comment. We treated the first feedback, if any, from buyer to seller and vice versa, that occurred within 60 days of the transaction, as feedback for that first transaction between partners.[6]

4.1 Sampling

The two-sided nature of the feedback systems we study poses particular challenges for data sampling. The unit of analysis in our maximum likelihood model is a transaction, and we assumed above that feedback strategy selection for all the transactions in the dataset are pairwise independent, conditional on the item price and feedback profiles. Yet suppose that each trader (buyer or seller) has an idiosyncratic individual propensity to choose one of the three feedback strategies (always give, never give, reciprocate). Then if a trader participates in multiple transactions, these transactions will be inter-dependent, thus violating the independence assumption.

We present below a method that yields consistent estimates in the face of this multiple-transaction, fixed-effect problem. Before we do, we explain why a couple of other seemingly natural methods do not work. One is to randomly select one from all transactions in which a given buyer participates, losing some observations but eliminating buyer interdependence. Unfortunately, since transactions involve trader pairs, this method will not, in general, eliminate all multiple transactions for

[5]Our transaction data began only in February 1999, while our feedback history goes back to the beginning of eBay. There is a chance that we included transactions that were not the first for a partnership, if the previous transactions were more than one month prior to our extraction window and no feedback had ever been given between the partners.

[6]It is possible that some partnerships conducted additional transactions beyond the first and provided feedback for subsequent transactions but not the first. Our data set would incorrectly attribute the first feedback with the first transaction. We believe that occurred infrequently and would introduce only random error, not systematic bias.

some sellers. Further, if multiple-transaction buyers are matched approximately randomly to sellers, then this sampling method will disproportionately eliminate altogether sellers who trade infrequently, biasing the resulting sample of sellers. Sampling both sides to eliminate multiple transactions for both buyers and sellers simply exacerbates the second problem.

Another approach is common: to specify a maximum likelihood model that explicitly accounts for the possibility of fixed effects, estimating them as nuisance parameters, to yield consistent estimates of the parameters of interest. However, in our sample of 959,657 items, there are 394,997 distinct buyers, and 133,697 distinct sellers. Thus, on average buyers in our sample appear in about two transactions, and sellers appear in about seven. Indeed, 50% of the traders had no more than two transactions, and 90% had no more than nine transactions. For most traders, this longitudinal dimension is much too small to rely on asymptotics for consistency. Greene (2004) found that with such a small number of repeated observations per agent, both fixed and random effects models produce results that are more biased than those from a pooled model in which trader-specific effects are ignored.

We now describe our method. We constructed two sub-samples, one a **buyer-unique** sub-sample containing a single randomly drawn transaction for each of the 394,997 unique buyers, and the other a **seller-unique** sub-sample containing a single randomly drawn transaction for each of the 133,697 unique sellers. We then estimate the model twice, once for each sub-sample. We obtain estimates of the parameters associated with buyer behavior from the buyer-unique sub-sample (discarding the estimates of the seller parameters), and we obtain seller parameter estimates from the seller-unique sub-sample.

Our sampling method ensures a representative sample. The buyer (seller) sub-sample is a representative cross-section sample of the buyer (seller) population. More importantly, the sampling method ensures unbiased parameter estimates for the uniquely-sampled side, despite potential biases in the parameter estimates for the other side. For example, consider the seller-unique sub-sample. For any fixed buyer parameters, the likelihood function is maximized at the same set of seller parameters.[7] Thus, because the seller-unique sub-sample has independence among seller transactions, we get an unbiased estimate of the seller parameters, even if the buyer parameter estimates are not accurate due to dependence among transactions involving the same buyer.

This method may not be efficient, compared to a hypothetical panel model, because we discard many observations. A panel model with fixed or random ef-

[7] One can easily verify this by taking the partial derivative of the log-likelihood function with respect to any seller parameter: no buyer parameters appear. This property is due to our assumption that in any transaction the buyer and the seller independently choose their feedback provision strategies.

fects, however, is not suited to our dataset, as argued above. Fortunately, with our rather large dataset, we can be somewhat profligate and still obtain rather precise estimates.

4.2 Model fit

We tested three different parametric functional forms, i.e., Lognormal, Gamma, and Weibull distributions, to estimate f_{sy} and f_{by}. The observed distribution and the estimated distributions using these three functional forms are graphed in Figure 2 for sellers and Figure 3 for buyers.

From a visual comparison of the estimated and the observed distributions, we believe that the lognormal model fits the observed distributions better than the other two. However, the results of Kolmogorov-Smirnov tests reject the null hypotheses that the observed and the predicted timing distributions are the same, for all three functional forms and for both timing distributions, i.e., f_{sy} and f_{by}, each at a statistically significant level (p-value < 0.001). This is understandable because with a large dataset (more than one hundred thousand observations), almost any deviation from a hypothesis will be statistically significant. Thus the Kolmogorov-Smirnov test results are overly precise for determining the reasonableness of the goodness of fit of our model.

We also conducted robustness test of our results using all three distribution functions for feedback timing (See appendix A for details). As varying the functional forms did not lead to qualitatively different results, we only report our results estimated using the lognormal distribution.

As a sanity check on the assumption that receiving feedback triggers reciprocation, Figure 4 and 5 plot the actual timing of feedback, for sellers and buyers respectively, after receipt of a feedback on day 15 and 35, as compared to the expected feedback if there were no reciprocators. Both show spikes in feedback giving the day or two immediately after receiving feedback. These spikes could be due, in part or in full, to a reminding effect. That is, people who would have given feedback anyway may do so earlier, reminded by the event of receiving feedback. Reciprocation effects cannot be read off simply from these graphs: we estimate the reciprocation effect through our model, by comparing the area under the curve to the right of the first feedback event to the expected area, rather than examining the shape of those curves. The existence of the spike, however, provides clear evidence that receipt of the first feedback has some effect on the other party's feedback actions, and thus it is reasonable to attribute modeled changes in the second player's actions to the effect of the first party's feedback.

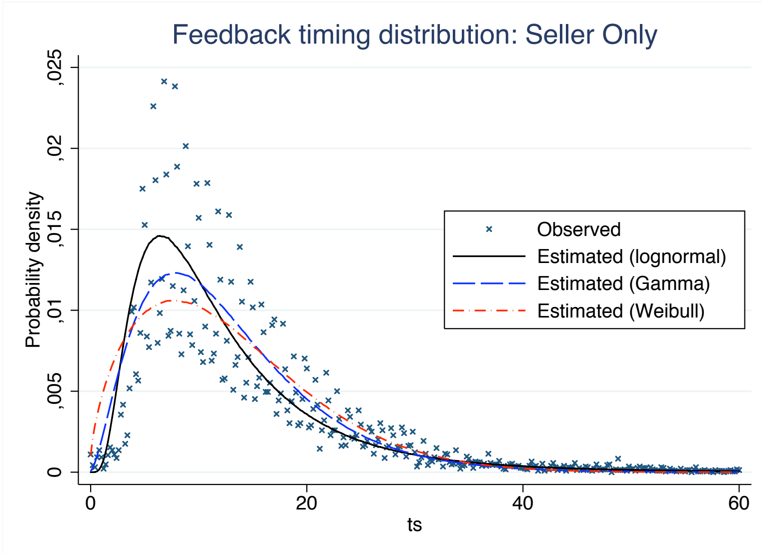

Figure 2: Sellers' feedback time distribution estimated using all the observations of Seller Only outcomes of our dataset.

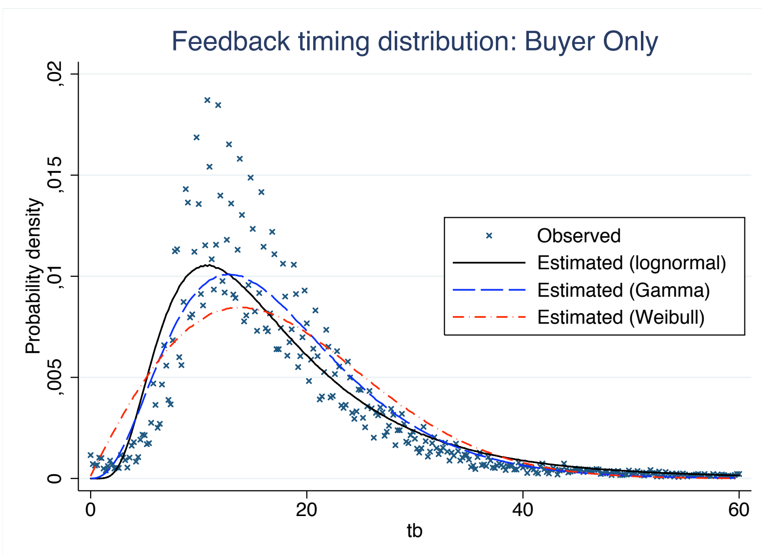

Figure 3: Buyers' feedback time distribution estimated using all the observations of Buyer Only outcomes of our dataset.

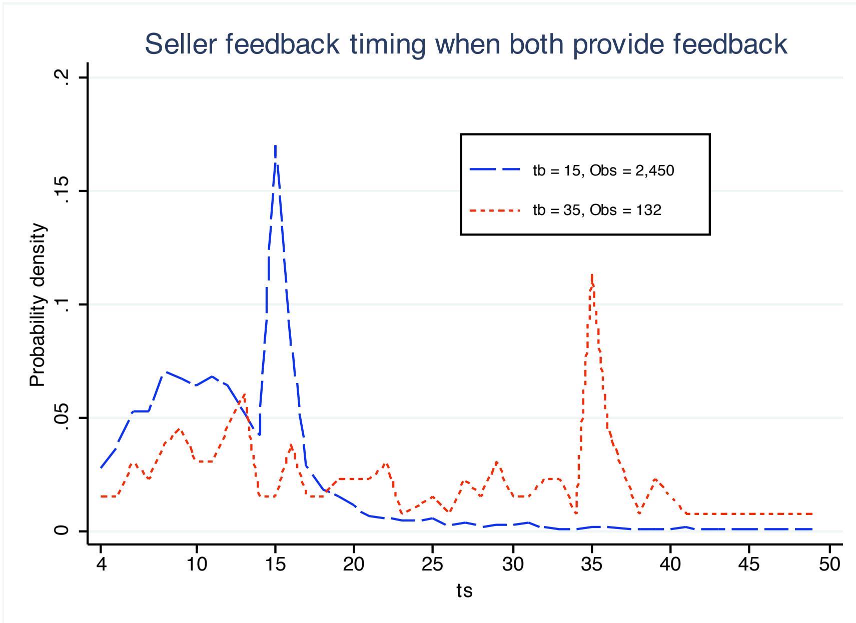

Figure 4: The distributions of seller feedback timing when buyers gave feedback on day 15 and day 35.

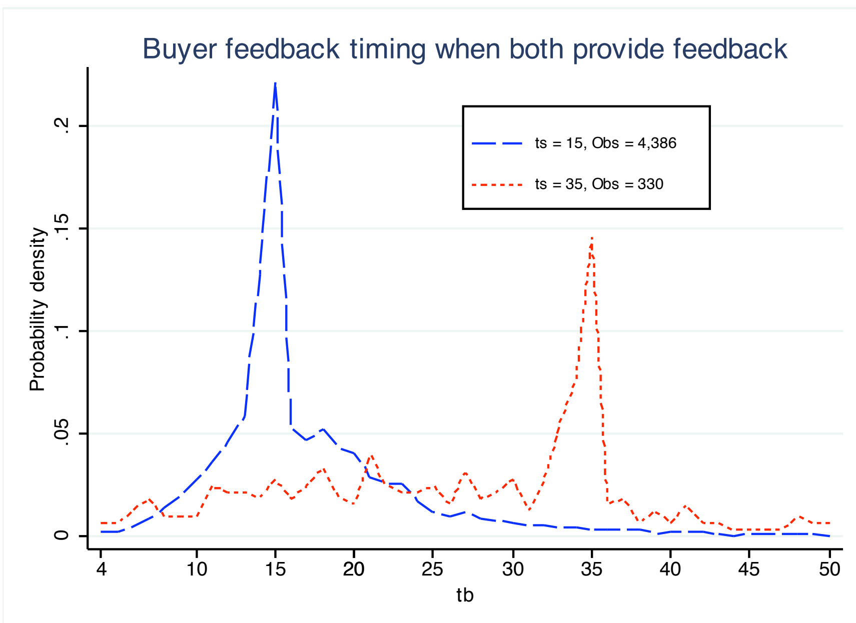

Figure 5: The distributions of buyer feedback timing when sellers gave feedback on day 15 and day 35.

	Buyer Unique		**Seller Unique**	
Outcomes	**# of Occurrences**	**%**	**# of Occurrences**	**%**
No Feedback	143,080	36.22%	45,927	34.35 %
Seller Only	73,276	18.55%	22,788	17.04 %
Buyer Only	40,669	10.3%	15,966	11.94 %
Seller First	90,564	22.92%	32,032	23.96 %
Buyer First	47,408	12%	16,984	12.71 %
Total	394,997	100%	133,697	100 %

Table 2: Distribution of feedback provision outcomes.

Roles	**Mean(Std. Err.)**	**Median**
Buyer	26(66)	6
Seller	78(155)	26

Table 3: Distribution of prior feedbacks received.

5 Empirical Results

Table 2 shows the distribution of feedback outcomes in both data sets: buyer-unique and seller-unique. Overall, about 35% of the transactions received no feedback; about 18% received feedback only from the seller; about 10% received feedback only from the buyer; and the remaining 35% received feedback from both parties. Table 3 shows the distribution of prior feedbacks received by the traders. Buyers have a mean of 26 feedbacks (with a standard error of 66), and a median of 6; sellers have a mean of 78 (with a standard error of 155), and a median of 26. Overall, sellers have more feedbacks than buyers. The distributions for both buyers and sellers are skewed: both distributions contain many traders with low feedback scores, and a few traders with high feedback scores. Note that when counting the total number of prior feedbacks for a trader, we followed eBay's practice (when it calculates feedback scores): we consider only the first feedback between a trading pair.

To study determinants of traders' choice of feedback provision strategies, we constructed six independent variables (see Table 4) to describe the observable context of the traders' transactions. We expect that the mix of feedback provision strategies will depend on the reputation profiles of the participants, since at various stages of a trader's reputation profile development, her knowledge of the system varies, as well as her needs to maximize gains from her reputation profile.

Variable name	**Definition**
newbuyer	1 if a buyer has received fewer than 5 feedbacks, and 0 otherwise
newseller	1 if a seller has received fewer than 5 feedbacks, and 0 otherwise
newint	an interaction term, equal to $newbuyer \times newseller$
$lnfbbuyerr$	the logarithm of the total number of feedback scores received by a buyer ($fbbuyerr$) plus one, thus $lnfbbuyerr = log(fbbuyerr + 1)$
$lnfbsellerr$	the logarithm of the total number of feedback scores received by a seller ($fbsellerr$) plus one, thus $lnfbsellerr = log(fbsellerr + 1)$
$lnprice$	$log(the\ sale\ price\ of\ the\ item)$

Table 4: Independent variables in the regression analyses.

To elaborate, we first classify traders into two distinct categories: new and experienced traders. New traders are those who have very little experience with the feedback system, reflected in the small number (e.g., $0 \sim 4$) of feedbacks they have previously received. After participating in a few more transactions and receiving more than five feedbacks, we classify them as "experienced". The dummy variables *newbuyer* and *newseller* take the value 1 if a buyer or a seller is new, and zero otherwise. [8] *newint* is an interaction term of *newbuyer* and *newseller* to separately identify cases in which both the seller and the buyer are new. We also constructed two continuous variables to proxy for the traders' experience, $lnfbbuyerr$ and $lnfbsellerr$. They are the logarithm of the total number of prior feedbacks received in the trader's life time with eBay, for buyers and sellers respectively.[9] Last, we allow the logarithm of the price of the item being sold, $lnprice$, to be a factor in both the buyer's and the seller's strategy choices.

[8]The choice of five as a threshold is arbitrary. We tested the sensitivity of our results to this threshold, and found that they *are* sensitive, but in the direction that reinforces our conclusion: newness matters, and experience effects show up after a modest number of feedbacks (see appendix B).

[9]To operationalize these two variables, we used $lnfbbuyerr = log(fbbuyerr + 1)$, and similarly for the seller $lnfbsellerr = log(fbsellerr + 1)$, to avoid the problem of an undefined logarithm of zero.

Dependent Var.	Independent Var.	Coef.	Std. Err.	z	P-value
B_y	**newbuyer**	-0.113	0.019	-6.040	0.000
	newseller	-0.282	0.024	-11.780	0.000
	newint	-0.042	0.028	-1.500	0.134
	lnfbbuyerr	0.194	0.005	37.540	0.000
	lnfbsellerr	-0.03	0.003	-8.380	0.000
	lnprice	0.045	0.003	13.480	0.000
	intercept	-0.480	0.025	-19.200	0.000
B_r	**newbuyer**	0.18	0.036	4.910	0.000
	newseller	-0.286	0.046	-6.160	0.000
	newint	0.01	0.054	0.160	0.871
	lnfbbuyerr	0.098	0.010	9.670	0.000
	lnfbsellerr	-0.069	0.006	-11.640	0.000
	lnprice	-0.047	0.006	-7.290	0.000
	intercept	-0.310	0.048	-6.480	0.000
S_y	**newbuyer**	-0.280	0.036	-7.890	0.000
	newseller	-0.086	0.034	-2.540	0.011
	newint	-0.117	0.040	-2.950	0.003
	lnfbbuyerr	0.076	0.007	10.450	0.000
	lnfbsellerr	0.201	0.007	27.610	0.000
	lnprice	-0.045	0.006	-8.020	0.000
	intercept	-0.268	0.043	-6.190	0.000
S_r	**newbuyer**	-0.308	0.08	-4.000	0.000
	newseller	0.494	0.069	7.130	0.000
	newint	0.089	0.086	1.030	0.303
	lnfbbuyerr	0.030	0.014	2.070	0.039
	lnfbsellerr	0.204	0.015	13.870	0.000
	lnprice	0.102	0.012	8.700	0.000
	intercept	-1.60	0.092	-17.260	0.000

Table 5: Estimated coefficients.

Role and Strategy	Probabilities at median	Probabilities at mean
B_y	0.38 (0.002)***	0.43 (0.002)***
B_r	0.23 (0.003)***	0.22 (0.002)***
B_n	0.39 (0.003)***	0.35 (0.002)***
S_y	0.47 (0.004)***	0.52 (0.002)***
S_r	0.20 (0.005)***	0.21 (0.003)***
S_n	0.32 (0.005)***	0.26 (0.003)***
Independent Var. Values		
	$newbuyer = 0$	$newbuyer = 0$
	$newseller = 0$	$newseller = 0$
	$newint = 0$	$newint = 0$
	$fbbuyerr = 6$ (median)	$fbbuyerr = 26$ (mean)
	$fbsellerr = 26$ (median)	$fbsellerr = 78$ (mean)
	$price = 45.43$ (mean)	$price = 45.43$ (mean)

Table 6: Estimated probabilities of each strategy being adopted by typical sellers and buyers.

Note: Asterisks indicate the level of statistical significance. *** means statistically significant at the 1% level, ** at 5% and * at 10%. The numbers in parentheses are the standard deviations.

In Table 5 we present the coefficients on these independent variables estimated using our maximum likelihood model for four dependent variables: B_y, B_r, S_y, and S_r (B_n and S_n can be derived from these four variables). In the subsequent columns, we list their standard errors, z values, and the p-values of two-sided tests that they are not different from zero. Most estimates of the coefficients are significantly different from zero at 1% level. The coefficients in this table are hard to interpret, however, as they do not easily translate into marginal effects. In the following sections we evaluate the marginal effects of the independent variables, contingent on scenarios in which the independent variables take on various values.

5.1 Distribution of feedback provision strategies

In Table 6 we report our estimates of the probabilities of buyers or sellers playing any of the three hypothesized strategies, evaluated at the median and the mean number of feedbacks for each trader type. The top half of the table contains the estimated probabilities and the bottom half of the table contains the values of the independent variables at which these probabilities are evaluated. All the estimated

Seller Strategy Prob.	Facing new buyer	Facing experienced buyer
$\widehat{S_y}$	$0.35(0.004)^{***}$	$0.45(0.004)^{***}$
$\widehat{S_r}$	$0.16(0.007)^{***}$	$0.18(0.005)^{***}$
$\widehat{S_n}$	$0.5(0.007)^{***}$	$0.38(0.005)^{***}$
S_y	$0.22(0.004)^{***}$	$0.33(0.004)^{***}$
S_r	$0.22(0.01)^{***}$	$0.24(0.006)^{***}$
S_n	$0.56(0.01)^{***}$	$0.43(0.005)^{***}$
$\widehat{S_y} - S_y$	$0.12(0.006)^{***}$	$0.12(0.005)^{***}$
$\widehat{S_r} - S_r$	$-0.06(0.012)^{***}$	$-0.06(0.007)^{***}$
$\widehat{S_n} - S_n$	$-0.06(0.012)^{***}$	$-0.06(0.007)^{***}$
Independent Var. Values		
	$newseller = 1$	$newseller = 1$
	$fbsellerr = 0$	$fbsellerr = 0$
	$\widehat{newseller} = 0$	$\widehat{newseller} = 0$
	$\widehat{fbsellerr} = 5$	$\widehat{fbsellerr} = 5$
	$newbuyer = 1$	$newbuyer = 0$
	$fbbuyerr = 0$	$fbbuyerr = 26$ (mean)
	$price = 45.43$ (mean)	$price = 45.43$ (mean)

Table 7: Seller behavior changes when new sellers become experienced.

probabilities shown in the table are statistically significantly different than zero at the 1% level. As we expected, a significant proportion of sellers and buyers are feedback reciprocators, and all three hypothesized strategies are being adopted for a substantial proportion of transactions. At the median levels (a buyer with a score of 6 buying a $45 item from a seller with score 26), there is a 38% probability the buyer will give feedback unconditionally and 39% probability she will abstain from giving feedback, with the remaining 23% probability she will be a reciprocator. On the other hand, at the median levels, 47% of sellers give feedback unconditionally; 32% abstain; and 20% reciprocate.

5.2 Strategy choices

Traders may use different feedback strategies at various stages of their career. They may also behave differently when facing different types of trading partners. In this section, we explore how traders' strategy choices vary based on the trading context.

Seller Strategy Prob.	**Facing new seller**	**Facing experienced seller**
$\widehat{B_y}$	$0.34(0.004)^{***}$	$0.37(0.002)^{***}$
$\widehat{B_r}$	$0.24(0.007)^{***}$	$0.22(0.003)^{***}$
$\widehat{B_n}$	$0.43(0.006)^{***}$	$0.41(0.003)^{***}$
B_y	$0.23(0.003)^{***}$	$0.27(0.002)^{***}$
B_r	$0.27(0.007)^{***}$	$0.26(0.003)^{***}$
B_n	$0.49(0.006)^{***}$	$0.47(0.002)^{***}$
$\widehat{B_y} - B_y$	$0.1(0.005)^{***}$	$0.1(0.003)^{***}$
$\widehat{B_r} - B_r$	$-0.04(0.009)^{***}$	$-0.04(0.004)^{***}$
$\widehat{B_n} - B_n$	$-0.06(0.009)^{***}$	$-0.06(0.003)^{***}$
Independent Var. Values		
	$newbuyer = 1$	$newbuyer = 1$
	$fbbuyerr = 0$	$fbbuyerr = 0$
	$\widehat{newbuyer} = 0$	$\widehat{newbuyer} = 0$
	$\widehat{fbbuyerr} = 5$	$\widehat{fbbuyerr} = 5$
	$newseller = 1$	$newseller = 0$
	$fbsellerr = 0$	$fbsellerr = 78$ (mean)
	$price = 45.43$ (mean)	$price = 45.43$ (mean)

Table 8: Buyer behavior changes when new buyers become experienced.

5.2.1 New traders become experienced

We compare new sellers' and experienced sellers' strategy choices in Table 7. The variables with hats indicate that they are about the experienced sellers, and those without hats are about new sellers. We consider these comparisons in two scenarios, one in which sellers face new buyers ($newbuyer = 1$, $fbbuyerr = 0$), and the other in which they face experienced buyers ($newbuyer = 0, fbbuyerr = 26$, where 26 is the mean of $fbbuyerr$). The results are quite similar in these two scenarios (compare across the two columns in Table 7), so we focus on one: sellers facing new buyers (the first column of results). In exactly the same format, Table 8 shows the comparisons of new buyers' and experienced buyers' strategy choices. Again we focus on the first column of the results in Table 8.

We found that the changes in behavior among sellers and buyers as they gain their first few feedbacks follow the same pattern: a reduction in the use of the "abstain" and "reciprocate" strategies, accompanied by a strong increase in "give unconditionally". The effects are larger in absolute and relative terms for sellers.

As new traders become experienced, they are more likely to give feedback unconditionally ($\widehat{S_y} - S_y = 0.12$, one-sided test: p-value < 0.001; and $\widehat{B_y} - B_y = 0.1$, one-sided test: p-value < 0.001), less likely to abstain from giving feedback ($\widehat{S_n} - S_n = -0.06$ one-sided test: p-value < 0.001; $\widehat{B_n} - B_n = -0.06$ (one-sided test: p-value < 0.001), and less likely to reciprocate ($\widehat{S_r} - S_r = -0.06$, one-sided test: p-value < 0.001; $\widehat{B_r} - B_r = -0.04$, one-sided test: p-value < 0.001). The probability of sellers (buyers) giving feedback unconditionally increases by 0.12 (0.1) point, which is a 55% (43%) increase from the new sellers' (buyers') probability of giving feedback unconditionally, $S_y = 0.22$ ($B_y = 0.23$), indicating that there is considerable learning going on among new traders during their first few trades. eBay provides FAQs and forums to facilitate learning. Also, the socially interactive nature of the feedback system makes it possible for new traders to learn by doing. Receiving feedback from one's trading partner informs or reminds the trader about the existence of the feedback system, and with it comes a link which guides her to return a feedback to her partner. In addition, receiving a feedback informs one about the social norm of feedback giving. Humans have the natural tendency to conform to social norms (Asch, 1956, Akerlof, 1980, Bernheim, 1994). All of these reasons point to the tendency that as new traders gain experience, they are more likely to be unconditional-givers, and less likely to abstain from giving feedback.

We do not have a strong conjecture as to whether new traders will be more likely to reciprocate as they become experienced. The results show that they will not. One possible explanation is that before new traders learned how to use the feedback system, some of them were "passive" reciprocators who did not know how to give feedback but were willing to reciprocate any feedback received. Some traders who had learned how to give feedback and also that others may reciprocate their feedback, started actively initiating feedback exchanges. Thus these reciprocators have instead become unconditional givers.

5.2.2 How experienced traders treat new traders

In our dataset, 180,433 (46%) out of the 394,997 buyers are new buyers, and 31,609 (24%) of the 133,697 sellers are new sellers. If the social nature of the feedback system plays a significant role in assisting new traders to learn, the speed at which they learn apparently depends on how soon they receive feedback. From the system designer's point of view, if the new traders received "special treatment" by veterans, they might learn faster. We have observed such special treatment in other communities, such as Wikipedia (Wikipedia, 2010). On eBay, do experienced traders also take up the responsibility of teaching new members? If so, they would be more likely to give feedback unconditionally to new traders than to other veterans.

Variable	Estimates
$S_y(\widehat{B})$	$0.50(0.004)^{***}$
$S_r(\widehat{B})$	$0.22(0.005)^{***}$
$S_n(\widehat{B})$	$0.28(0.005)^{***}$
$S_y(B)$	$0.43(0.004)^{***}$
$S_r(B)$	$0.20(0.007)^{***}$
$S_n(B)$	$0.37(0.006)^{***}$
$S_y(\widehat{B}) - S_y(B)$	$0.07(0.006)^{***}$
$S_r(\widehat{B}) - S_r(B)$	$0.02(0.009)^{**}$
$S_n(\widehat{B}) - S_n(B)$	$-0.08(0.007)^{***}$
Independent Var. Values	
	$newbuyer = 1$
	$fbbuyerr = 0$
	$\widehat{newbuyer} = 0$
	$\widehat{fbbuyerr} = 5$
	$newseller = 0$
	$fbsellerr = 78$ (mean)
	$price = 45.43$ (mean)

Table 9: Seller behavior changes when new buyers become experienced.

Table 9 and Table 10 contain our results on how traders behavior changes when their partners' reputation profile vary. Again, variables with hats indicate that they are for experienced traders. For instance, $S_y(\widehat{B})$ denotes the probability of a seller giving feedback unconditionally when facing an experienced buyer, and $S_r(B)$ denotes the probability of a seller playing the reciprocate strategy when facing a new buyer.

The results do not bear out our conjecture that there may be some "indoctrination" going on among the traders. We found experienced traders do not educate newbies by giving them more feedbacks; rather, they give newbies fewer. Both sellers and buyers are more likely to give feedback unconditionally to experienced traders than to newbies: $S_y(\widehat{B}) - S_y(B) = 0.07$ (one-sided test: p-value < 0.001) and $B_y(\widehat{S}) - B_y(S) = 0.04$ (one-sided test: p-value < 0.001); they are less likely to abstain when trading with experienced traders: $S_n(\widehat{B}) - S_n(B) = -0.08$ (one-sided test: p-value < 0.001) and $B_n(\widehat{S}) - B_n(S) = -0.05$ (one-sided test: p-value < 0.001). One possible explanation is that for a considerable proportion of experienced traders, the purpose of giving feedback unconditionally is to initiate feedback

Variable	**Estimates**
$B_y(\widehat{S})$	$0.43(0.003)^{***}$
$B_r(\widehat{S})$	$0.24(0.004)^{***}$
$B_n(\widehat{S})$	$0.33(0.003)^{***}$
$B_y(S)$	$0.41(0.003)^{***}$
$B_r(S)$	$0.22(0.004)^{***}$
$B_n(S)$	$0.38(0.003)^{***}$
$B_y(\widehat{S}) - B_y(S)$	$0.04(0.005)^{***}$
$B_r(\widehat{S}) - B_r(S)$	$0.007(0.007)$
$B_n(\widehat{S}) - B_n(S)$	$-0.05(0.006)^{***}$
Independent Var. Values	
	$newseller = 1$
	$fbsellerr = 0$
	$\widehat{newseller} = 0$
	$\widehat{fbsellerr} = 5$
	$newbuyerer = 0$
	$fbbuyererr = 26$ (mean)
	$price = 45.43$ (mean)

Table 10: Buyer behavior changes when new sellers become experienced.

exchanges. As new members are less familiar with the system, they may be less likely to "return the favor". As a result, experienced traders do not have sufficient incentives to give feedback to new traders.

Both veteran sellers and buyers are slightly more likely to reciprocate as their partners become experienced: $S_r(\widehat{B}) - S_r(B) = 0.02$ (one-sided test: p-value < 0.036) and $B_r(\widehat{S}) - B_r(S) = 0.007$ (not statistically significantly different from zero, but the change is in the right direction).

5.2.3 Gaining experience

Once an experienced seller has accumulated a substantial number of feedbacks, the marginal benefit of receiving a positive feedback diminishes, while the damage caused by a negative feedback becomes salient. To avoid receiving negative feedback, sellers may choose never to give feedback first, which means she would either reciprocate or abstain. In repeated interactions with the same buyers or in situations in which the seller's past feedback giving behavior with other buyers is observable, establishing a reputation as a reciprocator rather than an abstainer would help the seller create some reward and retaliation power. Taken together, we expect sellers with higher number of feedbacks to be more likely to reciprocate and less likely to give feedback unconditionally. For the buyers, we do not have a strong conjecture as to how they will behave differently as they gain experience.

In Table 11 we report the marginal effects on feedback giving evaluated at the sample medians and means of $fbbuyerr$ and $fbsellerr$. Each marginal effect is reported in terms of probability on the scale from 0 to 100 percentage points. For example, $fbsellerr\_on\_Sy$ — the effect of the total number of prior feedbacks received by the seller on her probability of playing strategy Y — is 0.11 percentage points with a standard deviation of 0.01, which reads: as a seller receives one more feedback, the probability that she gives feedback unconditionally increases by 0.11 percentage points, all else being equal.

As expected, we found sellers are more likely to reciprocate when they gain experience: $fbsellerr\_on\_Sr = 0.05$ (one-sided test: p-value < 0.001) at the median level and $fbsellerr\_on\_Sr = 0.01$ (one-sided test: p-value < 0.001) at the mean level. Thus, for a seller with a median (mean) number of feedbacks, adding 100 more feedbacks increases the probability that he reciprocates by 5 (1) percentage points. To our surprise, we found sellers are more likely to give feedback unconditionally as they gain experience ($fbsellerr\_on\_Sy = 0.11$, one-sided test: p-value < 0.001). That is, for every 100 feedbacks a seller receives, she is 11 percentage points more likely to give feedback unconditionally. This effect is smaller when evaluated at the mean level, but still significant: $fbsellerr\_on\_Sy = 0.03$, one-sided test: p-value < 0.001.[10] These increases in S_y and S_r come from a reduction in S_n: $fbsellerr\_on\_Sn = -0.16$ (or -0.05 evaluated at the means). Taken together, as experienced sellers receive even more feedback, they are more likely to give feedback unconditionally or reciprocate, and less likely to abstain from giving feedback.

Turning to the effect of increasing feedback on experienced buyers, we found buyers are more likely to give feedback unconditionally as they gain more experience: $fbbuyerr\_on\_By = 0.53$ at the median level (one-sided test: p-value < 0.001) and $fbbuyerr\_on\_By = 0.14$ at the mean level (one-sided test: p-value < 0.001). That is, buyers with 16 rather than 6 feedbacks go from 38% to 43.3% probability of choosing strategy Y of giving feedback unconditionally.[11] At the mean levels, buyers with 36 rather than 26 feedback go from 43% to 44.4% probability of choosing strategy Y.

[10] We do not mean to claim that the estimated marginal effect holds constant over a range of experience from 26 to 126 feedbacks; rather, we are simply rescaling the coefficient magnitude to improve comprehension of the effects.

[11] Again, this is not quite true, since the marginal effect does not hold constant over the range from 6 to 16 feedbacks.

Marginal effect	Evaluated at median	Evaluated at mean
$fbsellerr\_on\_Sy$	$0.11(0.01)^{***}$	$0.03(<0.00)^{***}$
$fbsellerr\_on\_Sr$	$0.05(0.01)^{***}$	$0.01(<0.00)^{***}$
$fbsellerr\_on\_Sn$	$-0.16(0.01)^{***}$	$-0.05(<0.00)^{***}$
$fbbuyerr\_on\_By$	$0.53(0.02)^{***}$	$0.14(<0.00)^{***}$
$fbbuyerr\_on\_Br$	$0.01(0.02)$	$-0.01(0.01)$
$fbbuyerr\_on\_Bn$	$-0.54(0.02)^{***}$	$-0.14(<0.00)^{***}$
$fbbuyerr\_on\_Sy$	$0.23(0.02)^{***}$	$0.06(0.01)^{***}$
$fbbuyerr\_on\_Sr$	$-0.03(0.03)$	$-0.01(0.01)^{*}$
$fbbuyerr\_on\_Sn$	$-0.19(0.03)^{***}$	$-0.04(0.01)^{***}$
$fbsellerr\_on\_By$	$-0.00(<0.00)$	$-0.00(<0.00)$
$fbsellerr\_on\_Br$	$-0.04(<0.00)^{***}$	$-0.01(<0.00)^{***}$
$fbsellerr\_on\_Bn$	$0.04(<0.00)^{***}$	$0.01(<0.00)^{***}$
$price\_on\_Sy$	$-0.05(<0.00)^{***}$	$-0.05(<0.00)^{***}$
$price\_on\_Sr$	$0.05(<0.00)^{***}$	$0.05(<0.00)^{***}$
$price\_on\_Sn$	$0.00(<0.00)$	$-0.00(<0.00)$
$price\_on\_By$	$0.03(<0.00)^{***}$	$0.03(<0.00)^{***}$
$price\_on\_Br$	$-0.03(<0.00)^{***}$	$-0.03(<0.00)^{***}$
$price\_on\_Bn$	$-0.01(<0.00)^{***}$	$-0.01(<0.00)^{***}$
Independent Var. Values		
	$newbuyer = 0$	$newbuyer = 0$
	$newseller = 0$	$newseller = 0$
	$fbbuyerr = 6$ (median)	$fbbuyerr = 26$ (mean)
	$fbsellerr = 26$ (median)	$fbsellerr = 78$ (mean)
	$price = 45.43$ (mean)	$price = 45.43$ (mean)

Table 11: Marginal effects on feedback giving evaluated at the sample medians and means of the number of feedbacks received by buyers and sellers.

5.2.4 How experienced traders are treated

We expect that experienced traders receive varying treatments based on the number of feedbacks they have accumulated. For example, as experienced buyers gather more (positive) feedback, they may seem more trustworthy as feedback givers. We expect that sellers are more willing to initiate a feedback exchange with trustworthy buyers.

As for experienced sellers, the more feedbacks they have, we would expect the lower the probability that they will receive feedback. Suppose seller A has 2000 feedbacks, while seller B has 100. Although both are experienced by our definition, we suspect that a buyer may be more likely to give feedback to B than to A, the reason being that the buyer feels her "vote" counts more for B than for A.

We did find that sellers are more likely to give feedback to buyers with more feedback: $fbbuyerr\_on\_Sy = 0.23$ (one-sided test: p-value < 0.001) at the median level, and $fbbuyerr\_on\_Sy = 0.06$ (one-sided test: p-value < 0.001) at the mean level. That is, for 10 new feedbacks a typical buyer receives, the probability that her partner seller gives feedback unconditionally to her increases by 2.3 percentage points at the median level, and by 0.6 percentage points at the mean level. This increase is mainly a decrease in S_n. Taken together, increasingly experienced buyers are more likely to always receive feedback from sellers. This result is consistent with our earlier results on how sellers treat buyers when we divide buyers into two groups: new versus experienced.

We did not find strong evidence to support our conjecture that buyers are less likely to give feedback unconditionally as sellers accumulate more experience: $fbsellerr\_on\_By = 0.00$. We did find (evaluating at the medians) that as sellers accumulate more feedback, buyers are less likely to reciprocate and more likely to abstain from giving them feedback: $fbsellerr\_on\_Br = -0.04$ (one-sided test: p-value < 0.001) and $fbsellerr\_on\_Bn = 0.04$ (one-sided test: p-value < 0.001).[12] It appears our conjecture is in the right direction: buyers stopped bothering about giving feedback (or returning feedback) to highly experienced sellers, but the effects are quite small.

5.2.5 Item value

Items with different values may spark different feedback behavior. We know that buyers pay more attention to their sellers' feedback profile when buying higher valued items (Ba and Pavlou, 2002). We suspect that buyers are more likely to give

[12] These results are consistent with Dellarocas and Wood (2008): buyers are more likely to give feedback to inexperienced sellers than to experienced sellers.

feedback unconditionally if the prices of the items are higher. Anticipating buyers' behavior, sellers may be safer to strategically reciprocate rather than give feedback first, to avoid negative feedback.

We did find that item value affects the strategy choices of both sellers and buyers. Buyers are more likely to give feedback unconditionally when the price of the item is higher: $price\_on\_By = 0.03$ (one-sided test: p-value < 0.001). We also found with high value items, buyers are less likely to reciprocate: $price\_on\_Br = -0.03$ (one-sided test: p-value < 0.001). Thus for a \$100 increase in the price of the item, the buyers are three percentage points more likely to give feedback unconditionally, and three percentage points less likely to reciprocate. On the seller side, we found sellers are more likely to reciprocate, and less likely to give feedback unconditionally with higher value items: $price\_on\_Sr = 0.05$ (one-sided test: p-value < 0.001), and $price\_on\_Sy = -0.05$ (one-sided test: p-value < 0.001). Thus with a \$100 increase in the item value, sellers are five percentage points more likely to be reciprocators, and five percentage points less likely to give feedback unconditionally. These results are consistent with findings from prior literature (Ba and Pavlou, 2002, Dellarocas and Wood, 2008): buyers are more likely to give feedback for higher value items, and they are pickier in assessing the quality of the services for higher value items; in response to buyers' behavior, sellers tend to strategically hold back their feedback to retain the option to give retaliating negative feedbacks.

6 Limitations

In our specification we assumed that buyers and sellers choose their feedback giving *strategies* independently, conditional on item price and feedback profiles. This does not mean that they select actions independently: in particular when someone chooses the reciprocation strategy, the action of giving or not giving feedback depends on whether the partner does. Independent strategy choice does mean that buyer and seller did not collude (e.g., by making an outside agreement) when selecting their strategies of whether to give feedback conditionally or unconditionally. This is a common assumption for online transactions, and is unlikely to be problematic in the case of most eBay transactions.

There is another implication of our strategy independence assumption that may be of greater concern. Our strategy independence assumption is an assumption that strategy selection is not conditional on variables unobserved by the econometrician, but observable (at least in part) by both parties.[13] We may expect that strategies in fact depend on such variables. For instance, a person's decision whether to

[13] We use "observable in part" to refer to the case in which the two parties may observe "different" variables, but these variables themselves have some common component, or a joint non-independent

send feedback might depend on the other party's timeliness in carrying out his or her part of the transaction.

The problem can be illustrated by referring back to one of the terms that enters our likelihood function; consider the likelihood given that the outcome of No Feedback is observed:

$$l(\theta; No\ Feedback) = (1 - S_y)(1 - B_y), \tag{3}$$

That the strategies S_y and B_y are chosen independently absent collusion is not problematic. However, real users — say, sellers — might choose a "give feedback" strategy S_y that depends on the quality of the buyer's performance (as in, "give feedback always, if the buyer sends a check within three days") or of communications between the buyer and seller (as in, "give feedback always, if the buyer announces that he will give feedback"), and if the buyer's strategy also depends on some of the same variables, then the actions chosen by S_y and B_y may be correlated, and our likelihood function is misspecified.

The problem of omitting correlated variables on which strategies are contingent is a fundamental identification problem for all latent variable models, and is not special to our dataset nor our specification. The empirical question is how good the conditional independence assumption is (that is, that the econometrician observes and conditions on all salient correlating variables). Studies of other datasets, with different or more conditioning variables, would be valuable to test the robustness of our main claims. All that we can say, as is usual, is that our results are conditional on the specification, which in our case means that we assume strategy choices by both buyer and seller are not (very much) conditional on unobservables.

Another limitation of our study pertains to our dataset. The transactions in our dataset occurred in 1999. eBay has revised its feedback rules multiple times over the last decade since our data collection.[14] A major change occurred in May 2008, when eBay removed sellers' ability to leave negative or neutral feedback on buyers, to free dissatisfied buyers from fear of retaliation when they leave negative feedback.[15] Under this new rule, sellers lost their power of retaliation, which was one of the reasons *why* sellers may have strategically reciprocated feedback during the time our data were collected. We do not know how such rule revisions may have

distribution. Then, the correlated part of these two variables can lead to correlation in the actions chosen.

[14]See details of these revisions at http://www2.ebay.com/aw/au/archive.shtml, retrieved on Sept 14, 2010.

[15]A number of measures were subsequently taken to protect sellers from buyers abusing their power conferred by this rule change, including enabling the buyer to revise her negative or neutral feedback if the seller managed to rectify the transaction problem. eBay also made a few other changes in May 2008, such as reducing the 90 day window of feedback giving to 60 days.

affected the population prevalence of eBay trader feedback provision strategies. Our methodological contribution provides a straightforward way to measure the effect of a rule change on feedback provision strategies using a before and after dataset.

In May 2007, eBay introduced Detailed Seller Ratings (DSRs) which enabled buyers to provide feedback on "four aspects of their transaction: accuracy of item description, communication, shipping time, and shipping and handling charges. The rating system is based on a one to five star scale, with one star being the lowest rating and five stars being the highest." The average ratings of a seller on all four aspects are displayed as part of her feedback profile.[16] The current system also displays the percentage of positive feedback out of the total number of positive and negative feedbacks received in the last 12 months.[17] These changes certainly enrich the display of the reputation profile, but the underlying mechanism remains largely the same. As long as the number of positive, neutral, and negative feedbacks are displayed, traders continue to care about receipt of these feedbacks when they formulate their feedback giving strategies.

eBay has also revised its feedback removal rules. At the time of our data collection, eBay did not allow revising feedbacks unless there was clear indication of feedback abuse. The current policy is that buyers can revise their negative or neutral feedback if the seller manages to rectify the problem that led to the feedback. Such a policy change might affect buyers' attitudes toward giving negative or neutral feedback. We did not study the content of feedback, just strategies for whether to provide feedback. As we discussed earlier, however, the opportunity to retaliate, and now to revise, may affect traders' choice of feedback provision strategies as well.

The evolution of the rules in eBay's feedback system reflects the fact that these rules are important: the system managers consider it important enough to have adjusted these rules multiple times to achieve better feedback outcomes. It also implies that traders do respond to these rules. Given the rule changes, some of our results may no longer hold for the current eBay feedback system. Nonetheless, our results continue to provide a baseline estimate of the prevalence of the three feedback giving strategies, for comparison with other feedback systems in other electronic markets, including today's eBay. And perhaps equally important, we have developed a straightforward method that can be used to identify the prevalence of feedback giving strategies in other data sets, under varying environments.

[16] See more details about DSRs at http://pages.ebay.com/help/feedback/detailed-seller-ratings.html, retrieved on Sept 14, 2010.

[17] When the percentage of positive feedback was first introduced in March 2003, it was calculated as the percentage of positives out of the total number of positive, neutral, and negative feedbacks. In July 2008, eBay removed neutrals from the calculation.

Because the feedback ecosystem is complex, no single study can account for the effects of future changes in the reputation system. Thus, having a reusable estimation method is valuable.

7 Conclusions

We developed an econometric model to study the feedback provision strategies used by participants in systems for bilateral interactions between strangers. We then applied our model to analyze the feedback provision strategies of eBay traders. We hypothesized that three types of feedback provision strategies were played by the traders: give (feedback) unconditionally, abstain (from giving feedback) unconditionally, and reciprocate. We found that all three types of strategies were being played by the traders. In particular, in transactions in which the buyer has the median number of feedbacks among all buyers, and the seller has the median number of feedbacks among all sellers, 38% of the buyers and 47% of the sellers give feedback unconditionally. This is quite high compared to theoretical predictions that the proportion of traders (buyers or sellers) who give feedback unconditionally would be minimal. We also found that in a substantial faction of cases, traders were strategic feedback reciprocators— 23% of the buyers and 20% of the sellers. We argue that the knowledge about the existence of these reciprocators may be a motivation for some traders to give feedback unconditionally, as they anticipate their partners to reciprocate. The remaining 39% of the time for buyers and 32% for sellers, the chosen strategy is not to provide feedback, regardless of whether the partner does.

eBay traders' feedback provision strategies evolve as they participate more in the marketplace. Both new buyers and new sellers become more likely to give feedback unconditionally after they experience their first few trades. Furthermore, as experienced traders continue to trade, they are also more likely to give feedback in general. Sellers are more likely to both reciprocate and give feedback unconditionally, and buyers are more likely to give feedback unconditionally. Overall, this is good news for eBay (and other trading systems that provide inter-partner performance feedback). As traders participate more in the system, they are more likely to be good citizens and hence provide feedback regardless of their trading partner's feedback giving actions.

Given our finding that new traders evolve into good citizens in terms of feedback giving, we expected there to be some "indoctrination" going on among eBay traders, but this does not appear to be true. Experienced sellers do not educate new buyers by giving them more feedback; neither do experienced buyers attempt to educate new sellers by giving them more feedback.

We also found that with high valued items, buyers are more likely to give feedback unconditionally but sellers are more likely to reciprocate. We speculate that as buyers care about the quality of high valued items more, they are more likely to pay attention and provide feedback. Experienced sellers would anticipate such behavior and strategically choose to reciprocate.

We also make a methodological contribution by building an econometric model to estimate feedback provision strategies in systems in which participants engage in bilateral interactions. Such types of systems can be electronic marketplaces, or systems that facilitate peer-to-peer sharing of services or resources. A special feature of our model is that the two parties' feedback provision strategies can be contingent on each other's actions, or not. Thus, either party can decide whether to give feedback based on what the other party does. With multinomial regressions, our model can be used to predict the participants' strategy choices based on their observable characteristics or the context of the interactions.

APPENDICES

A Robustness test of the timing distribution functional form assumptions

In section 5 we reported our results estimated under the assumption that feedback timing follows a lognormal distribution. To evaluate the robustness our results to the assumed functional form of the distribution, we also estimated the model with Weibull and Gamma distributions and report our results in Table 12.[18] For the lognormal function, we report the estimated coefficients of the independent variables on the four dependent variables, B_y, B_r, S_y, and S_r, and their associated p-values for two-sided tests that they are not different from zero. To compare the results estimated using Weibull and Gamma distributions to the lognormal distribution, we report the percentage change (under the column "% Δ") in each coefficient or parameter, that is, the percentage difference from the benchmark coefficient estimate under the lognormal assumption.

From Table 12 we conclude that the specific functional form assumption does not affect our results materially. Most percentage changes in the coefficients are under 10%, except for three (highlighted in bold), all of which are below 25%.

[18] Based on our priors, we only considered distributions with non-negative support, and which can be asymmetric with a long tail.

		Lognormal		Gamma	Weibull
		Coef.	**P-Value**	**% Err.**	**% Err.**
B_y	**newbuyer**	-0.11	-6.04	-0.48	-0.70
	newseller	-0.28	-11.78	-0.49	-0.20
	newint	-0.04	-1.5	0.43	2.32
	lnfbbuyerr	0.19	37.54	-0.16	-0.02
	lnfbsellerr	-0.03	-8.38	0.44	0.50
	lnprice	0.04	13.48	-1.56	-0.64
	intercept	-0.48	-19.2	0.94	0.31
B_r	**newbuyer**	0.18	4.91	2.75	0.41
	newseller	-0.29	-6.16	1.00	-1.14
	newint	0.01	0.16	**-20.25**	**-15.91**
	lnfbbuyerr	0.10	9.67	-0.66	-2.51
	lnfbsellerr	-0.07	-11.64	-0.85	-1.22
	lnprice	-0.05	-7.29	0.39	-0.44
	intercept	-0.31	-6.48	-4.69	0.46
S_y	**newbuyer**	-0.28	-7.89	0.32	1.15
	newseller	-0.09	-2.54	0.43	0.51
	newint	-0.12	-2.95	0.07	-0.68
	lnfbbuyerr	0.08	10.45	0.47	1.40
	lnfbsellerr	0.20	27.61	-0.04	-0.56
	lnprice	-0.04	-8.02	-0.03	2.24
	intercept	-0.27	-6.19	-0.52	-2.60
S_r	**newbuyer**	-0.31	-4	2.66	3.21
	newseller	0.49	7.13	-0.48	3.48
	newint	0.09	1.03	5.70	8.42
	lnfbbuyerr	0.03	2.07	**10.62**	7.43
	lnfbsellerr	0.20	13.87	0.25	0.85
	lnprice	0.10	8.7	-1.58	3.46
	intercept	-1.60	-17.26	-0.88	-4.44

Table 12: Coefficient robustness under different timing distribution assumptions. "% Δ" indicates the percentage difference from the benchmark coefficient estimate. Changes of more than 20% highlighted in bold.

B Robustness tests for the definition of "new trader"

We defined a "new" trader to be one with fewer than 5 feedbacks. In Table 13 we report the estimates obtained using different cutoffs, i.e., 3, 7, and 10 feedbacks. Most of the percentage errors are reasonably low, i.e., less than 50%. A few highlighted percentage errors in these are more than 50%. Some of the estimated coefficients that vary greatly across different cutoff values tend to have high p-values. For instance, the coefficient on *newint* for the dependent variable B_r has a percentage error of 1026.02% when estimated with Cutoff = 3. But the benchmark estimate of it with Cutoff = 5 has a p-value of 0.87. Thus this estimate was not statistically different from zero in the benchmark estimate. Examining other estimates with high percentage errors reveals that they do not alter our results qualitatively, though the size of the marginal effects might vary significantly depending on the cutoff values. This result reinforces our conclusion: "newness" matters.

C Simulation results

To validate our maximum likelihood model, we conducted Monte Carlo simulations. We report the simulation parameters and results in Table 14. Our simulated sample has 959,657 data points, the same sample size as our actual dataset. First, we arbitrarily picked a set of "true" parameter values, shown in the "True Coef." column, as the target true values to estimate. Next, we randomly generated the following independent variables each using a uniform distribution on $[0,1]$: *newbuyer*, $newseller$, *newint*, $lnfbbuyerr$, $lnfbsellerr$, and $lnprice$. Using the true coefficient values specified by us and the simulated independent variables, we generated the probabilities of strategy choices for each trader, i.e., either B_y and B_r or S_y and S_r depending on the role of the trader. We then use these probabilities to randomly select the "actual" strategy that the trader used. Last, using the selected strategies, and the parameters we specified for the timing functions, we generated a time stamp for each feedback given, if according to the trader's strategy she would give a feedback. Taking this simulated dataset, we then estimated our model with lognormal timing distributions, to obtain the coefficients shown in the "Coef." column. For each estimated coefficient, we also report its standard error (in the "SE" column), and the p-value of a two-sided test that the coefficient is not different from zero (in the "P-value" column). To analyze the simulation error, we compute the deviation measured as the number of standard deviations, and report this in the final column.

Overall, our results indicate that our model is valid. With approximately one million random draws, the distribution of simulation errors should converge close to a normal distribution. In Figure 6, we plot the cumulative distribution of the

		Cutoff = 5		Cutoff = 3	Cutoff = 7	Cutoff = 10
		Coef.	**P-Value**	**% Err.**	**% Err.**	**% Err.**
B_y	**newbuyer**	-0.11	0.00	**-110.86**	-19.35	-7.82
	newseller	-0.28	0.00	-7.54	24.36	13.77
	newint	-0.04	0.13	-34.93	-32.40	23.22
	lnfbbuyerr	0.19	0.00	-9.55	-2.26	-13.82
	lnfbsellerr	-0.03	0.00	14.16	8.32	-18.15
	lnprice	0.04	0.00	3.71	3.23	-1.21
	intercept	-0.48	0.00	8.23	2.40	28.01
B_r	**newbuyer**	0.18	0.00	-14.80	**-77.24**	**52.70**
	newseller	-0.29	0.00	9.38	19.16	-11.88
	newint	0.01	0.87	**1026.02**	**-258.55**	**392.86**
	lnfbbuyerr	0.10	0.00	-9.89	-37.60	23.65
	lnfbsellerr	-0.07	0.00	18.75	-5.65	-14.23
	lnprice	-0.05	0.00	-0.74	-0.35	18.68
	intercept	-0.31	0.00	-12.81	42.07	-18.19
S_y	**newbuyer**	-0.28	0.00	-10.18	27.41	11.87
	newseller	-0.09	0.01	**-222.20**	**57.76**	**209.76**
	newint	-0.12	0.00	-24.01	20.49	32.42
	lnfbbuyerr	0.08	0.00	12.62	14.32	-21.85
	lnfbsellerr	0.20	0.00	-12.09	3.88	4.12
	lnprice	-0.04	0.00	5.93	-0.43	-15.40
	intercept	-0.27	0.00	19.03	-26.86	23.56
S_r	**newbuyer**	-0.31	0.00	13.44	33.41	24.05
	newseller	0.49	0.00	-31.34	-20.20	15.88
	newint	0.09	0.30	5.40	**-60.22**	**-79.11**
	lnfbbuyerr	0.03	0.04	**62.42**	47.21	-39.76
	lnfbsellerr	0.20	0.00	-22.18	-5.95	9.00
	lnprice	0.10	0.00	0.72	-0.19	-2.87
	intercept	-1.60	0.00	6.61	-0.51	-1.26

Table 13: Robustness test result on the number of feedbacks that defines new trader. "% Δ" indicates the percentage difference from the benchmark coefficient estimate. Changes of more than 50% highlighted in bold.

Depd Var.	Indepd Var.	True Coef.	Coef.	SE	P-value	Err in SD
B_y	**newbuyer**	0.000	0.042	0.015	0.004	2.842
	newseller	-0.250	-0.256	0.015	0.000	-0.422
	newint	0.150	0.133	0.015	0.000	-1.140
	lnfbbuyerr	-0.300	-0.309	0.015	0.000	-0.589
	lnfbsellerr	0.510	0.525	0.015	0.000	1.023
	lnprice	-0.100	-0.106	0.015	0.000	-0.400
	intercept	0.300	0.288	0.018	0.000	-0.669
B_r	**newbuyer**	0.000	0.067	0.031	0.029	2.177
	newseller	-0.550	-0.548	0.031	0.000	0.057
	newint	0.520	0.501	0.031	0.000	-0.608
	lnfbbuyerr	-0.200	-0.226	0.031	0.000	-0.842
	lnfbsellerr	0.150	0.182	0.031	0.000	1.029
	lnprice	0.300	0.279	0.031	0.000	-0.682
	intercept	-0.300	-0.323	0.038	0.000	-0.621
S_y	**newbuyer**	0.000	0.024	0.014	0.079	1.756
	newseller	-0.300	-0.311	0.014	0.000	-0.822
	newint	-0.350	-0.345	0.014	0.000	0.330
	lnfbbuyerr	-0.500	-0.486	0.014	0.000	1.016
	lnfbsellerr	0.350	0.353	0.014	0.000	0.232
	lnprice	-0.350	-0.324	0.014	0.000	1.867
	intercept	-0.400	-0.437	0.017	0.000	-2.202
S_r	**newbuyer**	0.000	0.010	0.012	0.418	0.811
	newseller	-0.200	-0.198	0.012	0.000	0.149
	newint	0.800	0.800	0.013	0.000	0.016
	lnfbbuyerr	-0.800	-0.763	0.012	0.000	2.973
	lnfbsellerr	0.450	0.423	0.012	0.000	-2.149
	lnprice	0.100	0.119	0.012	0.000	1.515
	intercept	0.400	0.372	0.016	0.000	-1.797
Timing Parameters	μ_{sy}	3.500	3.496	0.007	0.000	-0.542
	σ_{sy}	2.000	1.997	0.005	0.000	-0.698
	μ_{by}	4.000	4.000	0.003	0.000	0.033
	σ_{by}	2.000	2.001	0.002	0.000	0.487

Table 14: Simulation Results with a sample of 959,657 data points.

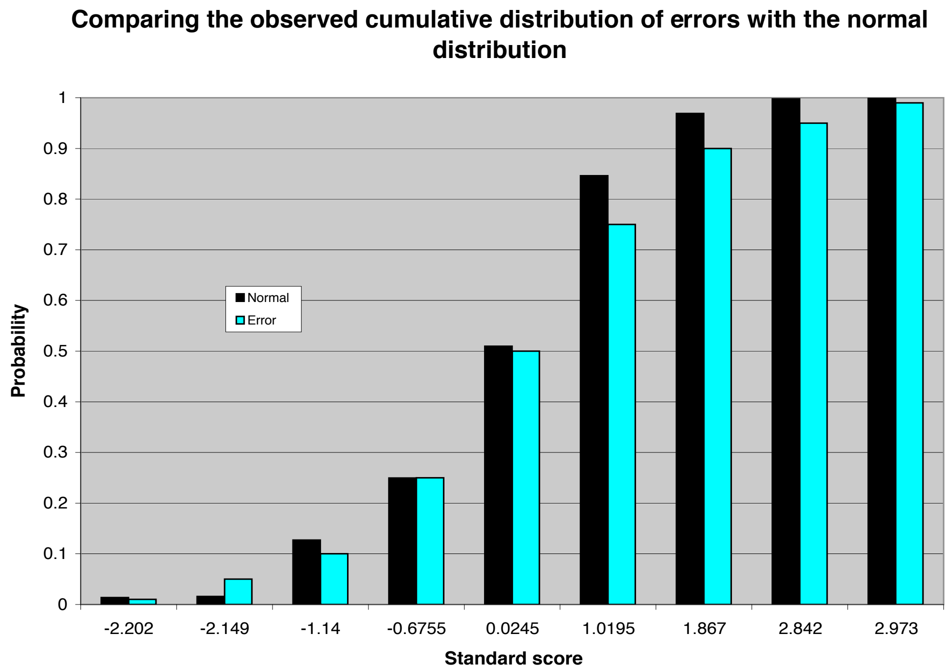

Figure 6: Comparing the cumulative distribution of the simulation errors with the normal distribution.

Overall, our results indicate that our model is valid. With approximately one million random draws, the distribution of simulation errors should converge close to a normal distribution. In Figure 6, we plot the cumulative distribution of the normalized simulation errors against the predicted normal error distribution. The tails are a bit heavier, but overall the fit is quite good and there are no outliers.

References

Akerlof, G. A. (1980): "A theory of social custom, of which unemployment may be one consequence," *The Quarterly Journal of Economics*, 94, 749–775.

Andreoni, J. and J. H. Miller (1993): "Rational cooperation in the finitely repeated prisoner's dilemma: Experimental evidence," *The Economic Journal*, 103, 570–585.

Asch, S. E. (1956): "Studies of independence and conformity: A minority of one against a unanimous majority," *Psychological Monographs*, 70.

Ba, S. and P. A. Pavlou (2002): "Evidence of the effect of trust building technology in electronic markets: Price premiums and buyer behavior," *MIS Quarterly*, 26, 243–268.

Bernheim, B. D. (1994): "A theory of conformity," *The Journal of Political Economy*, 102, 841–877.

Bolton, G., B. Greiner, and A. Ockenfels (2009): "Engineering trust - reciprocity in the production of reputation information," Working paper.

Cabral, L. and A. Hortaçsu (2010): "The dynamics of seller reputation: Theory and evidence from eBay," *Journal of Industrial Economics*, 58, 54–78.

Dellarocas, C. (2003): "The digitization of word-of-mouth: promise and challenges of online reputation mechanisms," *Management Science*, 49, 1407–1424.

Dellarocas, C. and C. A. Wood (2008): "The sound of silence in online feedback: Estimating trading risks in the presence of reporting bias," *Management Science*, 54.

Fehr, E. and S. Gächter (2002): "Altruistic punishment in humans," *Nature*, 415.

Gächter, S. and A. Falk (2002): "Reputation and reciprocity: Consequences for the labour relation," *Scandinavian Journal of Economics*, 104, 1–26.

Gazzale, R. (2004): *Giving Gossips Their Due: Information Provision in Games with Private Monitoring*, Ph.D. thesis, University of Michigan, Ann Arbor, MI.

Greene, W. (2004): "The behaviour of the maximum likelihood estimator of limited dependent variable models in the presence of fixed effects," *Econometrics Journal*, 7, 98–119.

Jian, L. and J. K. MacKie-Mason (2008): "Why share in peer-to-peer networks?" in *ICEC '08: Proceedings of the 10th international conference on Electronic commerce*, New York, NY, USA: ACM, 1–8.

Khopkar, T., X. Li, and P. Resnick (2005): "Self-selection, slipping, salvaging, slacking, and stoning: the impacts of negative feedback at eBay," in *EC '05: Proceedings of the 6th ACM conference on Electronic commerce*, New York, NY, USA: ACM, 223–231.

Resnick, P. and R. Zeckhauser (2002): "Trust among strangers in Internet transactions: Empirical analysis of eBay's reputation system," in M. R. Baye, ed., *The Economics of the Internet and E-Commerce*, *Advances in Applied Microeconomics*, volume 11, Amsterdam: Elsevier Science, 127–157.

Resnick, P., R. Zeckhauser, E. Friedman, and K. Kuwabara (2000): "Reputation systems," *Communications of the ACM*, 43, 45–48.

Samuelson, P. A. (1954): "The pure theory of public expenditure," *Review of Economics and Statistics*, 36, 387–389.

Wikipedia (2010): "Wikipedia:Welcoming committee/Welcome to Wikipedia," See http://en.wikipedia.org/wiki/Wikipedia:Welcoming_committee, retrieved on Sep 14, 2010.