# THE SOUND OF SILENCE IN ONLINE FEEDBACK: ESTIMATING TRADING RISKS IN THE PRESENCE OF REPORTING BIAS

CHRYSANTHOS DELLAROCAS

CHARLES A. WOOD

Abstract. Most online feedback mechanisms rely on voluntary reporting of privately observed outcomes. This introduces the potential for reporting bias, a situation where traders exhibit different propensities to report different outcome types to the system. Unless properly accounted for, reporting bias may severely distort the distribution of *public* feedback relative to the underlying distribution of *private* transaction outcomes and, thus, hamper the reliability of feedback mechanisms. This study offers a method that allows users of feedback mechanisms where both partners of a bilateral exchange are allowed to report their satisfaction to "see through" the distortions introduced by reporting bias and derive unbiased estimates of the underlying distribution of privately observed outcomes. A key aspect of our method lies in extracting information from the number of transactions where one or both trading partners choose to remain silent. We apply our method to a large data set of eBay feedback. Our results confirm the widespread belief that eBay traders are more likely to post feedback when satisfied than when dissatisfied. Furthermore, we provide rigorous evidence for the presence of positive and negative reciprocation among eBay traders. Most importantly, our analysis derives unbiased estimates of the risks that are associated with trading on eBay that, we believe, are more realistic than those suggested by a naïve interpretation of the unusually high (> 99%) levels of positive feedback currently found on that system.

## 1. Introduction

Online feedback mechanisms have become an important component of electronic business, helping to elicit good behavior and cooperation among loosely connected and geographically dispersed economic agents (Dellarocas 2003). For example, eBay's feedback mechanism is the primary means through which eBay elicits honest behavior and, thus, facilitates transactions among strangers over the Internet (Resnick and Zeckhauser 2002).

Since most details of commercial transactions are privately observed by the parties involved, the majority of online feedback mechanisms rely on voluntary self-reporting of transaction outcomes. As a consequence, not every transaction receives feedback. More importantly, self-reporting opens

*Date*: January 2006; Revised: June 2006, October 2006.

the door to several forms of *reporting bias*: traders may selectively choose to report certain types of outcomes and not others. If reporting bias is severe enough, public feedback provides a distorted view of the risks that are associated with trading in a given market. Its usefulness, both in deterring fraud and in informing buyers, then becomes severely diminished.

There are important indications that reporting bias is present in online feedback: Feedback in most systems is overwhelmingly positive. For example, more than 99% of all feedback posted on eBay is positive (Resnick and Zeckhauser, 2002; Kauffman and Wood, 2005). A naïve reading of this empirical fact may lead one to conclude that more than 99% of eBay transactions result in satisfactory outcomes. Such a conclusion runs against widespread reports of consumer fraud in online auctions. For example, Internet Auctions accounted for 16% of all consumer fraud complaints received by the Federal Trade Commission in 2004, the highest level of fraud of any Internet transaction type (see http://www.consumer.gov/sentinel/).

One possible explanation for this discrepancy is that, whereas satisfied traders generally report their satisfaction online, dissatisfied traders often prefer to remain silent. The reciprocal nature of auction feedback is considered by many as the main reason behind such reporting bias. Specifically, it is widely believed (though, so far, not rigorously proven) that many traders choose to remain silent because they are afraid that, if they report their negative experience, their partner will "retaliate" by posting negative feedback for them as well.

The presence of reporting bias in online feedback mechanisms has been mentioned by several authors (Reichling 2004; Klein et al. 2005; Hu et al. 2006). However, so far there has not been an attempt to quantify the degree to which it is present in a given system or an assessment of the extent to which it distorts the distribution of published feedback relative to the underlying distribution of transaction outcomes that traders privately experience.

We fill this gap by offering what we believe is the first quantitative method that can assess and repair the impact of reporting bias on feedback mechanisms. Given a sufficiently large sample of online feedback our method derives quantitative estimates of user propensities to report various types of outcomes to the system. Based on these estimates, the method then derives unbiased estimates of the distribution of *private* transaction outcomes that is most likely to have produced the target sample of *public* online feedback. Our approach, thus, enables traders to "see through" potentially biased online feedback and, thus, to obtain a more reliable picture of the risks associated with transacting in a given space. The method is fairly general and applies to a wide variety of *bidirectional feedback mechanisms*, that is, mechanisms that allow both partners of a bilateral exchange to rate each other.

We apply our method to a large data set of online feedback obtained from eBay. Our results provide rigorous evidence supporting the fact that eBay traders are more likely to report satisfactory outcomes than mildly unsatisfactory outcomes. (Reporting probabilities go up again when traders are very dissatisfied.) Furthermore, we show that a trader's propensity to post feedback is highly sensitive to her partner's reporting actions. In addition to confirming that unfavorable feedback increases the other trader's propensity to post unfavorable feedback in return, we show that favorable feedback increases the other trader's propensity to post favorable feedback in return (when satisfied) and to withhold posting unfavorable feedback (when mildly dissatisfied). Overall, our results indicate that reciprocity is an important driver of reporting behavior on eBay.

Our method is able to disentangle a trader's reporting behavior from the transaction outcome she has observed and, thus, to derive estimates of the distribution of *private* outcomes that is most likely to have produced the *public* feedback patterns observed in our data set. Our most detailed model estimates that, on average, eBay buyers walk away from a transaction satisfied 78.9% of the time, mildly dissatisfied 20.4% of the time and very dissatisfied 0.7% of the time. The corresponding estimates for sellers are 85.7%, 13.7% and 0.6% respectively. These, we believe, are more realistic estimates of trader satisfaction rates than the 99% rate suggested by a naïve interpretation of the percentage of positive feedback currently found on eBay.

An important element of our method consists in extracting information from the temporal order of buyer and seller feedback submission, as well as from the fraction of transactions where one or both traders choose to remain silent. Our work, therefore, demonstrates that a trader's choice to not post feedback provides important information that can be exploited to "see through" the distortions introduced by reporting bias. eBay currently does not publish any information on the number of a trader's transactions where the partner did not provide feedback. We argue that this omission makes it difficult for eBay traders to accurately assess the risks that are inherent in trading online and diminishes the effectiveness of its feedback mechanism.

We believe that this work contributes on several fronts. First, we offer a general methodology that can be applied to assess the presence of reporting bias and the distribution of privately observed transaction outcomes in a wide variety of bidirectional feedback mechanisms. Second, our analysis is the first to derive quantitative estimates of average trader satisfaction and feedback reporting behavior on eBay. Third, our results suggest that the impact of reciprocity in people's online reporting behavior is more complex than previously thought: Whereas we confirm prior conjectures suggesting that the fear of retaliation might discourage some traders from reporting bad outcomes, we also find evidence suggesting that the expectation of *positive* reciprocation may be partly responsible for the high levels of feedback contribution on eBay. Last, but not least, we demonstrate how one can extract useful information from a trader's decision to *not* post feedback and argue that the number of *silent transactions* (i.e., transactions for which no feedback was posted) should become a standard part of a trader's feedback profile on eBay and other feedback mechanisms.

The rest of the paper is structured as follows. Section 2 describes our data set. Section 3 introduces a family of models that draw inferences from the relative frequency of different types of feedback observed in a sufficiently large sample of transactions. Section 4 extends our baseline models to take into consideration the temporal order of feedback submission; we show that this extension allows the analyst to derive more precise estimates of reporting bias, including estimates of how one partner's feedback affects the other partner's subsequent reporting actions. Section 5 further extends our modeling technology to derive transaction-specific estimates of trader satisfaction and reporting behavior. Finally, Section 6 discusses the managerial implications of this work and lists opportunities for future research.

## 2. Data Set

Our data set consists of 51,062 rare coin auctions that took place from April 24, 2002 to September 11, 2002 on eBay. These auctions include items from 6,242 distinct sellers and 16,405 distinct buyers. We only consider auctions that resulted in a transaction (i.e., auctions that received at least one bid and where the secret reserve price, if it exists, was met). Our data set includes auction information (auction id, item description, ending time, selling price, number of bids), seller information (eBay id, seller feedback profile information), and winning bidder information (eBay id, buyer feedback profile information). In addition, our data set contains full information (date and time of feedback, auction id, rater's eBay id, feedback type: positive, neutral, or negative, associated text comment) related to feedback posted for these auctions by both buyers and sellers within a 90-day window following the closing time of the corresponding auction.[1]

[1]eBay encourages traders to leave feedback within 90 days after the termination of an auction and does not guarantee that traders will be able to leave feedback after that period. Empirical evidence suggests that feedback left after 90 days is extremely rare.

| | |
|---|---|
| Total number of auctions | 51,062 |
| Distinct buyers | 16,045 |
| Distinct sellers | 6,242 |

| | Min | Mean | Median | Max |
|---|---|---|---|---|
| Buyer feedback score | -2 | 63.6 | 38 | 814 |
| Seller feedback score | -1 | 154.8 | 101 | 852 |
| Bids per auction | 1 | 5.7 | 4 | 64 |
| Auction closing price | \$0.01 | \$52.98 | \$15.50 | \$16,500 |

| | Number of Auctions | % of Total |
|---|---|---|
| Auctions where seller left comment | 39,561 | 77.48% |
| Auctions where buyer left comment | 34,614 | 67.79% |
| Auctions where both left comment | 29,139 | 57.07% |
| Auctions where none left comment | 6,026 | 11.80% |
| Auctions where seller commented first | 30,524 | 59.78% |
| Auctions where buyer commented first | 14,902 | 29.18% |

Table 1. Key descriptive statistics of our data set.

| | | Number of Auctions | % of Total |
|---|---|---|---|
| Auctions where seller left comment | | 39,561 | |
| | positive | 39,275 | 99.28% |
| | neutral | 54 | 0.14% |
| | negative | 232 | 0.58% |
| Auctions where buyer left comment | | 34,614 | |
| | positive | 34,260 | 98.98% |
| | neutral | 163 | 0.47% |
| | negative | 191 | 0.55% |

Table 2. Breakdown of posted feedback into positive, neutral, and negative.

Table 1 summarizes some key descriptive statistics of our data set. All metrics of feedback score refer to the "standard" eBay feedback score.[2] We observe that feedback contribution is substantial: 77% of auctions receive a comment from the seller and 67% of auctions a comment from the buyer. Sellers post the first comment almost twice as often as buyers, reflecting the fact that the outcome (good/bad) of a transaction typically becomes clear to the seller sooner than to the buyer.[3]

Table 2 breaks down posted feedback into positive, neutral, and negative comments. The breakdown is consistent with that reported by most other studies, exhibiting an overwhelming (99%) preponderance of positive feedback. Our data set can be further divided into *feedback patterns*

[2]The "standard" eBay feedback score is equal to the sum of positive ratings minus the sum of negative ratings posted on behalf of a trader by distinct partners over the course of that trader's entire "career" on eBay. In the event that a trader receives multiple ratings from the same partner, all of them count as one. See Resnick and Zeckhauser (2002) for a detailed description of eBay's feedback mechanism.

[3]In most cases, a transaction is settled for the seller as soon as he receives money from the buyer. The buyer, on the other hand, must receive and examine the goods before she can determine her level of satisfaction.

| Buyer Comment Type | Seller Comment Type | Who comments first? | Number of auctions | % of Total |
|---|---|---|---|---|
| + | + | b | 9303 | 18.22% |
| + | + | s | 19613 | 38.41% |
| + | 0 | b | 7 | 0.01% |
| + | 0 | s | 3 | 0.01% |
| + | - | b | 12 | 0.02% |
| + | - | s | 4 | 0.01% |
| 0 | + | b | 18 | 0.04% |
| 0 | + | s | 57 | 0.11% |
| 0 | 0 | b | 7 | 0.01% |
| 0 | 0 | s | 0 | 0.00% |
| 0 | - | b | 5 | 0.01% |
| 0 | - | s | 2 | 0.00% |
| - | + | b | 4 | 0.01% |
| - | + | s | 60 | 0.12% |
| - | 0 | b | 1 | 0.00% |
| - | 0 | s | 0 | 0.00% |
| - | - | b | 31 | 0.06% |
| - | - | s | 12 | 0.02% |
| + | S | b | 5318 | 10.42% |
| 0 | S | b | 64 | 0.13% |
| - | S | b | 93 | 0.18% |
| S | + | s | 10220 | 20.02% |
| S | 0 | s | 39 | 0.08% |
| S | - | s | 163 | 0.32% |
| S | S | | 6026 | 11.80% |
| | | *Total* | *51062* | *100%* |

**Legend**

| | |
|---|---|
| *Comment Types* | |
| + | Positive feedback |
| 0 | Neutral feedback |
| - | Negative feedback |
| S | No feedback (silence) |
| *Who comments first?* | |
| b | Buyer |
| s | Seller |

Table 3. Consideration of the type and relative order of comments posted by the buyer and seller for a given transaction gives rise to 25 distinct feedback patterns.

according to the type and temporal order of feedback posted by buyers and sellers in their respective auctions. For example, one pattern consists of auctions where the seller posts positive feedback first, and the buyer responds with positive feedback. Another pattern consists of auctions where the buyer posts positive feedback and the seller remains silent. If we consider all possible combinations of each trader's feedback behavior (positive, neutral, negative feedback plus silence) and all possible temporal orderings of comments (buyer rates first, seller rates first) we obtain 25 mutually disjoint feedback patterns, including a pattern that contains auctions where both the buyer and the seller remain silent. Table 3 lists all 25 feedback patterns and their relative incidence in our data set. Our subsequent analysis of reporting bias and private transaction outcomes is heavily based on the extraction of information from the relative incidence of these 25 patterns in online feedback data.

Figure 1 plots the empirical distributions of feedback posting times relative to the corresponding auction's closing time. We observe that positive feedback is posted relatively soon and that buyer feedback lags seller feedback by 3-4 days. This lag is intuitive, since sellers are typically in a position to post positive feedback for a buyer as soon as they receive payment, whereas buyers

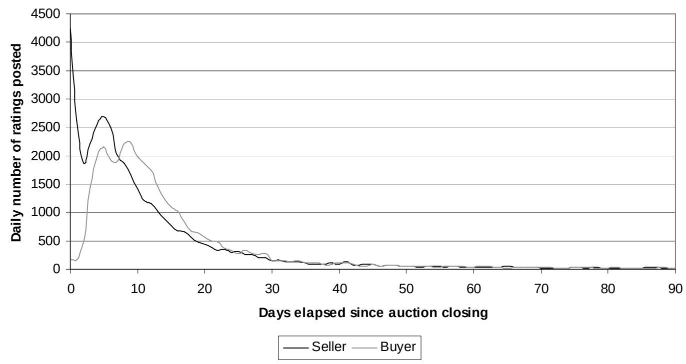


(a) Positive feedback

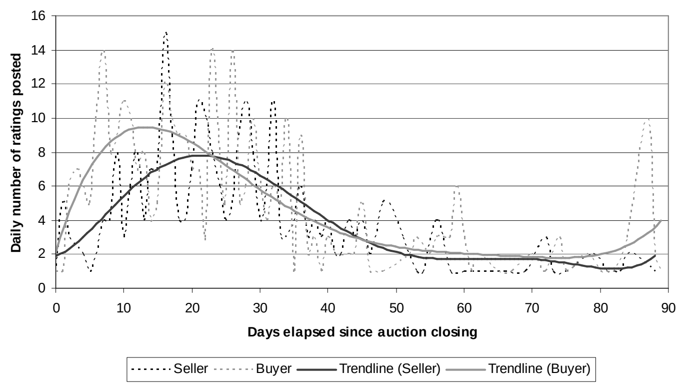


(b) Neutral and negative feedback

FIGURE 1. Empirical distribution of feedback posting times relative to the corresponding auction's closing time.

need to receive and examine the goods before they determine their level of satisfaction. We also observe that both the seller and buyer feedback distributions are bimodal; furthermore the locations of the corresponding modes are highly correlated. We hypothesize that this bimodality is a simple consequence of the variety of payment types supported by eBay: some buyers pay by credit card and Paypal, and can, therefore immediately communicate their payment to the seller. Others use checks that take a while to reach the seller. The seller, typically will not post feedback until payment is received and confirmed. We, therefore, hypothesize that the second mode in seller feedback corresponds to transactions where payment is by check. Furthermore, the seller will only ship the goods after he confirms receipt of payment. This leads to a corresponding delay in buyer feedback for such transactions.

Since unfavorable (negative or neutral) feedback is far less common, the shapes of the corresponding empirical distributions (dotted lines) are less regular; we, therefore, supplement them by sixth-degree polynomial trendlines (solid lines). We observe that unfavorable feedback is posted later than positive feedback. This makes sense since problematic transactions are usually associated with payment and/or shipment delays and additional communication between buyer and seller before traders "give up" on each other and post unfavorable feedback.

We conclude this section with a brief look at the text comments associated with eBay feedback. We read all text comments in our data set seeking insights related to the types of trading risks that are present on eBay. We found that the majority of positive comments do not contain very useful information for our purposes.[4] Neutral and negative comments, on the other hand, were quite insightful, as they point to a number of different sources of trader dissatisfaction. Table 4 groups neutral and negative comments found in our data set into a number of "problem areas." According to our results, the most common source of buyer complaints were transactions where the promised items were never received (40% of negatives; 7% of neutral), followed by items whose quality was inferior to what was expected (35% of negatives; 50% of neutral). Slow shipping or other post-sale communication problems with the seller accounted for around 20% of complaints. Interestingly, about 5% of buyer complaints referred to situations where the seller backed out of the transaction after the auction's completion. Finally, a few buyers complain because they found the shipping charges to be too high relative to the actual postage paid by the seller. The majority (81%) of negative seller comments relate to bidders who back out of their commitment to buy the items they won. Poor communication and unreasonable post-sale buyer demands accounted for another 13% of seller complaints. Finally, 2.6% of seller complaints refer to buyers who are slow in sending payment to the seller.

The differences in the distribution of neutral and negative comments among problem areas suggest that, in line with eBay's suggested guidelines, traders are more likely to post neutral feedback when problems are mild (e.g. unhappy with item quality, poor communication), reserving negative feedback for situations where problems are severe (e.g. item not received, buyer never sent payment).

## 3. Baseline models

Our baseline models draw inferences from the relative frequencies of different types of feedback observed in a sufficiently large sample of transactions. We show that, if we can assume the absence of

[4]It is customary for eBay traders posting positive comments to use excessive praise and colorful language. The most common text comment associated with positive feedback reads something like "Great transaction - A+++." More colorful comments, like "Faster than a cheetah chasing an antelope on an african plain!!!", referring to one of the authors' very prompt payment for an auction he once won, are not uncommon.

| Buyer Comments | Negative Comments | | Neutral Comments | |
|---|---|---|---|---|
| | Number | % | Number | % |
| Item not received | 76 | 39.79% | 12 | 7.36% |
| Unhappy with item quality | 66 | 34.55% | 82 | 50.31% |
| Slow shipping, poor communication | 30 | 15.71% | 38 | 23.31% |
| Seller backed out of transaction | 9 | 4.71% | 9 | 5.52% |
| Shipping charges deemed excessive | 5 | 2.62% | 7 | 4.29% |
| Other * | 5 | 2.62% | 15 | 9.20% |
| *Total unfavorable comments:* | *191* | *100%* | *163* | *100%* |

| Seller comments | Negative Comments | | Neutral Comments | |
|---|---|---|---|---|
| | Number | % | Number | % |
| Buyer never sent payment | 189 | 81.47% | 26 | 48.15% |
| Poor communication, unreasonable buyer | 31 | 13.36% | 16 | 29.63% |
| Slow payment | 6 | 2.59% | 3 | 5.56% |
| Other * | 6 | 2.59% | 9 | 16.67% |
| *Total unfavorable comments:* | *232* | *100%* | *54* | *100%* |

* No clear reason given.

TABLE 4. Summary of transaction problems associated with neutral and negative comments in our data set.

strategic misreporting, such models are always identifiable in mechanisms that allow both partners to rate one another.

### 3.1. **Basic concepts and identifiability results.**

We refer to feedback mechanisms that allow both partners of a transaction to rate each other as *bidirectional* feedback mechanisms. Consider a bidirectional feedback mechanism that allows traders to self-report privately observed transaction outcomes. Assume that each transaction can result in one out of $N \geq 2$ distinct outcome types (good, average, bad, etc.) for *each* partner. A transaction's outcome need not be identical for both partners. For example, a transaction where the buyer promptly sends payment and the seller ships back damaged items would leave the seller satisfied but the buyer dissatisfied. When both partners' satisfaction levels are taken into consideration, each transaction can, thus, have $N^2$ distinct outcomes.

Each of the two partners is given the option to report her level of satisfaction by posting one out of $M \geq 1$ available *feedback types* (e.g. integers between 1 and 5) on a public website. Each partner also has the right to remain silent. In the rest of the section it will be convenient to treat silence as an additional feedback type. If we treat silence as feedback, our mechanism supports $M+1$ distinct feedback types per partner. When both partners rate this gives rise to $(M+1)^2$ distinct feedback patterns per transaction, including patterns where one or both traders post no feedback.

The public website aggregates posted feedback and publishes the relative frequencies of all distinct feedback patterns. An important objective of this paper is to explore under what conditions one can estimate the incidence probabilities of the $N^2$ *private* transaction outcomes from the relative frequencies of the $(M+1)^2$ *public* feedback patterns.

The problem can be cast as a latent variable problem (Bollen 1989). Let $i_b, i_s \in \{1, ..., N\}$ denote the outcome experienced by a transaction's buyer and seller respectively. Similarly, let $j_b, j_s \in \{1, .., M+1\}$ denote the type of feedback posted by each partner (including no feedback). Let $\pi_{i_b i_s}$ denote the probabilities of the $N^2$ private transaction outcomes and let $\rho^k_{j_k|i_k}$ denote the probability that trader $k$ ($k = b(uyer)$ or $s(eller)$) reports feedback type $j_k$ conditional on having observed outcome $i_k$. The following system of polynomial equations then relates the probability $F_{j_b j_s}$ of observing feedback pattern $j_b j_s$ to the unknown probabilities $\pi_{i_b i_s}$ and $\rho^k_{j_k|i_k}$:

$$F_{j_b j_s} = \sum_{i_b=1}^{N} \sum_{i_s=1}^{N} \pi_{i_b i_s} \rho^b_{j_b|i_b} \rho^s_{j_s|i_s} \qquad j_b, j_s \in \{1, .., M+1\} \tag{1}$$

Recall (Bollen 1989) that a simultaneous equation model is identifiable if and only if it satisfies the *order* condition (number of independent equations $\geq$ number of independent unknowns) and the *rank* condition (rank of Jacobian matrix equal to number of independent unknowns). The above model has $(M+1)^2-1$ independent equations (system (1) consists of $(M+1)^2$ equations that satisfy $\sum_{j_b=1}^{M+1} \sum_{j_s=1}^{M+1} F_{j_b j_s} = 1$). Unknowns include $N^2 - 1$ independent private outcome probabilities (there are $N^2$ unknown outcome probabilities $\pi_{i_b i_s}$ that must satisfy $\sum_{i_b=1}^{N} \sum_{i_s=1}^{N} \pi_{i_b i_s} = 1$) and $MN$ independent reporting probabilities for *each* of the two partners (for each partner $k$ and each of the $N$ possible outcomes, there are $M+1$ unknown reporting probabilities $\rho^k_{j|i}$ that must satisfy $\sum_{j=1}^{M+1} \rho^k_{j|i} = 1$). The total number of independent unknowns is, thus, $N^2 - 1 + 2MN$. Unless $M$ is substantially larger than $N$, elementary algebra shows that the number of unknowns is greater than the number of equations. However, even when $M$ is sufficiently large so that the order condition holds, the following proposition shows that model (1) fails to uniquely identify the unknown outcome and reporting probabilities.

**Proposition 1.** *Model (1) is not identifiable for any $M \geq 1$, $N \geq 2$.*

Things get better if we can draw upon domain knowledge to reduce the number of unknown reporting probabilities $\rho^k_{j|i}$. For example, suppose that we can assume that some reporting probabilities are equal to zero (because, for example, we know that traders never post unfavorable feedback when happy or favorable feedback when unhappy). A special case of practical interest is one where we can assume that:

(A1) there is a one-to-one mapping between transaction outcomes and report types

(A2) traders either truthfully report the transaction outcome they observe or remain silent.

The following result then holds:

**Proposition 2.** *Under assumptions (A1) and (A2) model (1) is identifiable for all* $N \geq 2$.

3.2. **Application to eBay.** This section applies the previous results in the context of our eBay data set. Recall that eBay's feedback mechanism supports three distinct feedback types (positive, neutral, negative). To take advantage of Proposition 2 we develop a model where:

- each transaction has three possible outcomes (good, mediocre, bad) for each trader: *good* outcomes imply that the trader's expectations were met or surpassed, *mediocre* outcomes imply mild dissatisfaction and *bad* outcomes imply severe dissatisfaction.
- there is a one-to-one mapping between transaction outcomes and feedback types (good↔positive, mediocre↔neutral, bad↔negative)
- each trader either truthfully reports the feedback type that corresponds to the outcome she observed, or stays silent

The assumption that eBay traders are generally honest when posting feedback can be justified by the preponderance of one-time transactions in large-scale electronic markets. If a trader is unlikely to transact with the same partner again in the future, elementary game theory predicts that she is indifferent between truthful and untruthful reporting (Dellarocas 2005). Therefore, although we do not preclude isolated incidents of false reporting, we assume that *systematic* strategic misreporting does not take place at the population level.

Appendix I lists the equations derived from specializing model (1) to the context of the above domain ($M = N = 3$). In the rest of the paper we will refer to this model as *Model A*. The model has $(M+1)^2 = 16$ equations and $N^2 + 2M = 15$ unknowns. The notation $F_{j_b j_s}$ denotes the probability of observing an eBay transaction where the buyer posts feedback $j_b \in \{+, 0, -, S\}$ and the seller posts feedback $j_s \in \{+, 0, -, S\}$. Symbols $+, 0, -$ denote positive, neutral and negative feedback respectively; $S$ denotes "silence". For example, $F_{+S}$ denotes the probability of observing a transaction where the buyer posts positive feedback whereas the seller remains silent. The 15 unknown model parameters are also listed in Appendix I. Outcome probabilities are denoted $\pi_{i_b i_s}$:

the first subscript denotes the outcome $i_b \in \{G, M, B\}$ privately observed by the buyer and the second the outcome $i_s \in \{G, M, B\}$ privately observed by the seller. Symbols $G, M, B$ denote a $G$ood, $M$ediocre and $B$ad outcome respectively. Since we have assumed honest reporting and a one-to-one mapping between outcome and feedback types, reporting probabilities are simply denoted $\rho_i^k$; $k$ indicates the type of trader ($b$uyer or $s$eller) and $i \in \{G, M, B\}$ indicates the outcome observed by that trader (the type of feedback posted is implied by the type of outcome observed). For example, $\rho_G^b$ denotes the probability that a buyer who observed a good outcome will post (positive) feedback; accordingly, $1 - \rho_G^b$ is the probability that a buyer who observed a good outcome will remain silent.

The key to understanding the form of our model's equations is the observation that, given our assumption of truthful reporting, feedback patterns where both partners post feedback reveal the underlying latent outcome experienced by both partners. In contrast, feedback patterns where one of the partners stays silent only reveal the latent outcome experienced by the vocal partner. All 16 Model A equations are straightforward consequences of this observation.

Model A parameters can be estimated using the maximum likelihood method. Specifically, we can treat the manifest feedback pattern $j_b j_s$ of each transaction as a random variable that follows a multinomial distribution with 16 possible outcomes, whose respective occurrence probabilities $F_{j_b j_s}$ are given by the model's equations. The unknown parameter estimates $\pi_{i_b i_s}, \rho_i^k$ then are the ones that maximize the corresponding log-likelihood function:

$$\mathcal{L} = \sum_{j_b, j_s \in \{+,0,-,S\}} N_{j_b j_s} log\left(F_{j_b j_s}\right) \tag{2}$$

subject to the constraints:

$$\pi_{i_b i_s}, \rho_i^k \in [0,1] \quad \text{and} \quad \sum_{i_b, i_s} \pi_{i_b i_s} = 1 \tag{3}$$

$N_{j_b j_s}$ denotes the number of transactions in our data where we observe feedback pattern $j_b j_s$ and $F_{j_b j_s}$ is the right hand side of the corresponding Model A equation. It is known (see, for example, Greene 2002, p. 127) that, for sufficiently large samples, the maximum likelihood method leads to (asymptotically) consistent, minimum variance unbiased estimators.

| Parameter | ML Estimates | | Confidence Intervals | | |
|---|---|---|---|---|---|
| | Mean | Std. Error | 2.50% | Median | 97.50% |
| $\pi_G^b = \pi_{GG} + \pi_{GM} + \pi_{GB}$ | 0.815 | 0.014 | 0.785 | 0.815 | 0.842 |
| $\pi_M^b = \pi_{MG} + \pi_{MM} + \pi_{MB}$ | 0.174 | 0.015 | 0.148 | 0.174 | 0.205 |
| $\pi_B^b = \pi_{BG} + \pi_{BM} + \pi_{BB}$ | 0.011 | 0.003 | 0.005 | 0.010 | 0.018 |
| $\pi_G^s = \pi_{GG} + \pi_{MG} + \pi_{BG}$ | 0.886 | 0.009 | 0.865 | 0.887 | 0.900 |
| $\pi_M^s = \pi_{GM} + \pi_{MM} + \pi_{BM}$ | 0.104 | 0.009 | 0.089 | 0.103 | 0.126 |
| $\pi_B^s = \pi_{GB} + \pi_{MB} + \pi_{BB}$ | 0.010 | 0.003 | 0.006 | 0.010 | 0.016 |
| $\rho_G^b$ | 0.824 | 0.015 | 0.798 | 0.823 | 0.855 |
| $\rho_M^b$ | 0.017 | 0.002 | 0.014 | 0.017 | 0.022 |
| $\rho_B^b$ | 0.411 | 0.139 | 0.223 | 0.381 | 0.771 |
| $\rho_G^s$ | 0.868 | 0.009 | 0.855 | 0.867 | 0.890 |
| $\rho_M^s$ | 0.011 | 0.002 | 0.008 | 0.011 | 0.015 |
| $\rho_B^s$ | 0.476 | 0.128 | 0.282 | 0.453 | 0.795 |

Table 5. Maximum likelikood estimates of Model A parameters.

Table 5 lists the parameter estimates that solve the above constrained maximization problem. To facilitate the interpretation of the results we do not show the estimates of all nine joint outcome probabilities $\pi_{i_b i_s}$ but instead list the estimates of the *marginal* trader satisfaction probabilities

$$\pi_i^b = \sum_j \pi_{ij} \qquad \pi_i^s = \sum_j \pi_{ji}$$

i.e. the marginal probabilities that the buyer and seller will observe a good, mediocre and bad outcome respectively.

The rest of the section discusses the insights provided by the above results.

*Trader satisfaction.* The results of Table 5 imply that, on average, eBay transactions leave buyers satisfied 81.5% of the time, mildly dissatisfied 17.4% of the time and very dissatisfied 1.1% of the time. The corresponding figures for sellers are 88.6%, 10.4% and 1% respectively. These figures are in line with common sense and more credible than the 99% satisfaction rate that is tacitly implied by the percentage of positive feedback on eBay. Note that a buyer satisfaction rate of 81.5% does not imply that eBay sellers behave badly 18.5% of the time. As discussed in Section 2, buyer dissatisfaction has many sources, some of which are based on the buyer's own misunderstanding of the way eBay works. A similar disclaimer applies to seller dissatisfaction. Note, also, that the confidence intervals of our trader satisfaction probability estimates $\pi_G^b$ , $\pi_G^s$ are relatively broad

(78% to 84% for buyers; 86% to 90% for sellers). As we will show in the next section, tighter intervals can be obtained by considering the timing of feedback submission.

*Reporting behavior.* Our results show that satisfied eBay traders exhibit a relatively high propensity to post positive feedback (82% for buyers; 87% for sellers). Mildly dissatisfied traders, instead, prefer to remain silent, posting neutral feedback only about 1% of the time. The reporting probabilities of very dissatisfied traders are substantially higher (41% for buyers; 47% for sellers), but still lower than those of satisfied traders. Notice that the confidence intervals of $\rho_B^b$ and $\rho_B^s$ are very broad, suggesting that our data does not contain sufficient information to allow precise inferences with respect to the reporting behavior of very dissatisfied traders. Importantly, however, the 95% confidence intervals of $\rho_G^k$, $\rho_M^k$ and $\rho_B^k$ are mutually disjoint (for both buyers and sellers). This allows us to conclude (with 95% confidence) that, on average, eBay traders are more likely to report good outcomes than bad outcomes and more likely to report bad outcomes than mediocre outcomes.

## 4. Feedback Timing and Reciprocity

Our baseline model (Model A) offers valuable initial insight into the trading risks and feedback reporting behavior of eBay traders. On the positive side, the model is easy to understand and relatively straightforward to estimate. On the negative side, Model A ignores important aspects of the eBay domain. Specifically, our model assumes that each trader's reporting behavior is independent of his partner's reporting behavior. Such an assumption might be plausible in feedback mechanisms that simultaneously publish both traders' ratings for each other; it is less plausible on eBay, where traders post feedback asynchronously and where posted feedback is immediately visible to the other trader. In such a setting it is very likely that the feedback posted by the trader who reports first will have an impact on the subsequent reporting behavior of the other trader.

A long literature in psychology and behavioral economics offers powerful evidence for the importance of reciprocity in human interactions (Fehr and Gachter 2000). Reciprocal behavior is the pattern of behavior where people respond to friendly or hostile actions with similar actions. In the context of this work, reciprocity arguments suggest that receipt of positive feedback from a transaction partner might make a trader (who has not yet posted any feedback) more likely to report a good outcome and to withhold reporting a mediocre or bad outcome (*positive reciprocation*). Similarly, receipt of neutral or negative feedback might make the same trader more likely to report a mediocre or bad outcome and to withhold reporting a good outcome (*negative reciprocation*).

Understanding the extent to which reciprocity affects traders' online reporting behavior is an interesting question in its own right. It complements our understanding of reporting bias in feedback mechanisms and has important implications for their design (see Section 6). At the same time, failure to recognize that the first mover's feedback might affect the other trader's reporting behavior could lead to inaccurate estimates of the underlying private outcome probabilities.[5]

Motivated by the preceding arguments this section extends our baseline model to capture the potential for reciprocity-driven changes in the second mover's reporting behavior.

4.1. **A model with feedback timing.** Attempts to incorporate the impact of reciprocity into our baseline model quickly run into identifiability problems. The first step towards building such a model is to refine the patterns of observable feedback that drive our original model to specify not only the type of feedback posted by each trader but also which trader (buyer, seller) posts the first feedback. This increases the number of distinct observable patterns from 16 to 25 (see Table 3). The second step is to, similarly, refine the equations of Model A, replacing each of the first 9 equations (that describe patterns where both traders post feedback) with two equations, one where the buyer rates first and one where the seller rates first. The third step is to condition the second mover's reporting probabilities on the type of rating $j$ posted by his trading partner. Given 3 possible outcomes $i$, 3 possible partner ratings $j$ and 2 trader types $k$, this step introduces 18 additional model parameters $\tilde{\rho}^k_{i|j}$. We end with a model that has 25 independent equations and 33 unknowns (the original 15 unknowns plus the 18 conditional reporting probabilities). In the rest of the paper, we will refer to this model as *Model B.* Model B's equations and unknown parameters are listed in Appendix II.

Model B fails to satisfy the order condition and, thus, is not identifiable. However, it serves as a stepping stone for constructing an identifiable model. As we will show, identification can be obtained if we extend Model B to take into consideration the *time* at which each trader posts her respective feedback (relative to the beginning of the corresponding transaction).

The basic assumption that underlies our new model is that a trader's *time-to-feedback* correlates with the type of private outcome observed by that trader. This assumption can be justified if traders post feedback soon after they determine the outcome of a transaction and if the time it takes to determine a transaction's outcome correlates with the outcome type. The last assumption is plausible on eBay since good outcomes are likely to be determined sooner than mediocre or bad

[5]The terms *first mover* and *second mover* refer to the relative order of feedback submission.

outcomes. A good outcome is one where the buyer promptly sends payment to the seller and the seller promtly ships the promised goods to the buyer. Mediocre and bad outcomes, in contrast, are characterized by payment and/or shipment delays, unsatisfactory goods that are often returned to the seller and additional communication between buyer and seller as they try to resolve the dispute. The worse the final outcome, the longer it usually takes before the situation settles.

We now develop expressions that model the probability of observing a randomly chosen trader post feedback of a given type at a given point in time. Our expressions aim to capture the resulting population-level behavior and are independent of whether any or all traders behave strategically or not. We distinguish between the case where a trader rates before his partner and the case where a trader rates after his partner.

*Trader rates before partner.* As long as the partner has not yet posted feedback, we assume that a trader's time-to-feedback (conditional on the trader having decided to leave feedback) is governed by a failure time distribution $z_i^k(t) \equiv z(t; \theta_i^k)$ that depends on the type of trader $k$ (buyer, seller) and the type of outcome $i$ observed by that trader; $z(\cdot; \theta_i^k)$ denotes a suitable parametric family (e.g. Lognormal, Weibull, Gamma, etc.) whose parameter vector $\theta_i^k$ is a function of the privately observed outcome $i$ and the trader's type $k$; $Z(\cdot; \theta_i^k)$ denotes the corresponding CDF.

Since traders might decide to stay silent, the density function $r_i^k(t)$ (CDF $R_i^k(t)$) that characterizes the probability of observing a trader of type $k$ who has experienced outcome $i$ post feedback *(before his partner)* at time $t$, must also include the trader's reporting probability $\rho_i^k$:

$$r_i^k(t) = \rho_i^k z(t; \theta_i^k) \qquad R_i^k(t) = \rho_i^k Z(t; \theta_i^k) \tag{4}$$

*Trader rates after partner.* If a partner posts feedback at time $t_0$, her action is likely to affect the trader's *subsequent* conditional probability of posting feedback for the same transaction, given that he hasn't done so already. At the population level the latter quantity is simply the *hazard rate* of $r_i^k(t)$. Accordingly, we model the impact of partner feedback on the trader's subsequent propensity to rate by assuming that the partner's action multiplies the *hazard rate* of the trader's density function $r_i^k(t)$ by a factor $\alpha_{i|j}^k$ for all $t \geq t_0$. Factor $\alpha_{i|j}^k$ has the following interpretation: If $\alpha_{i|j}^k$ is greater than 1, this implies that *on average* partner feedback $j$ *increases* trader $k$'s subsequent propensity to report outcome $i$. If, on the other hand, factor $\alpha_{i|j}^k$ is less than 1, then partner feedback $j$ *decreases* trader $k$'s subsequent propensity to report outcome $i$.

The following proposition provides the analytic form of a failure time distribution whose original hazard rate gets multiplied by a constant factor $\alpha$ at all times $t \geq t_0$:

**Proposition 3.** *Let $f(t)$, $F(t)$ denote a failure time density and its CDF respectively. At time $t_0$, an external shock multiplies the hazard rate of $f(t)$ by a factor $\alpha$ for all $t \geq t_0$. The probability $\widetilde{f}(t|t_0)$ (CDF $\widetilde{F}(t|t_0)$) of observing the pertinent event occurring at time $t \geq t_0$ is then given by:*

$$\widetilde{f}(t|t_0) = \alpha f(t) \left( \frac{1 - F(t)}{1 - F(t_0)} \right)^{\alpha - 1} \qquad \widetilde{F}(t|t_0) = 1 - (1 - F(t_0)) \left( \frac{1 - F(t)}{1 - F(t_0)} \right)^{\alpha} \tag{5}$$

The density function $\widetilde{r}^k_{i|j}(t|t_0)$ (CDF $\widetilde{R}^k_{i|j}(t|t_0)$) that characterizes the probability of observing a trader post feedback after his partner is an immediate corollary of Proposition 3.

**Corollary 1.** *Let $r^k_i(t)$ ($R^k_i(t)$) denote the probability (CDF) of observing a randomly chosen trader of type $k$ who has experienced outcome $i$ post feedback at time $t$ if his trading partner has not yet posted feedback. Assume that the partner posts feedback $j$ at time $t_0$. The pdf (CDF) that describes the probability of observing a randomly chosen trader $k$ post feedback at time $t \geq t_0$ is then equal to:*

$$\widetilde{r}^k_{i|j}(t|t_0) = \alpha^k_{i|j} r^k_i(t) \left( \frac{1 - R^k_i(t)}{1 - R^k_i(t_0)} \right)^{\alpha^k_{i|j} - 1} \qquad \widetilde{R}^k_{i|j}(t|t_0) = 1 - (1 - R^k_i(t_0)) \left( \frac{1 - R^k_i(t)}{1 - R^k_i(t_0)} \right)^{\alpha^k_{i|j}} \tag{6}$$

We have now derived expressions of the probability of observing a randomly-chosen trader post feedback at time $t$ conditional on his type, observed outcome and his partner's reporting action up to that point. Going back to our (underidentified) Model B, if we replace all $\rho^k_i$ with $r^k_i(t)$ and all $\tilde{\rho}^k_{i|j}$with $\widetilde{r}^k_{i|j}(t|t_0)$, we obtain a new model that incorporates the times of feedback submission. We will refer to this new model as *Model C* (see Appendix III). Each of Model C's 25 equations describes the density function $f^k_{j_b j_s}(t_b, t_s)$ of observing a feedback pattern where trader $k$ posts feedback first, the buyer posts $j_b$ at time $t_b$ and the seller posts $j_s$ at time $t_s$. For the most part, Model C's equations are straightforward extensions of the corresponding equations of Model B. The only area where some explanation is needed are equations that involve silent traders (the model's last 7 equations). Model C assumes the existence of a data set that contains observations of feedback posted from the beginning of each transaction up until some *cutoff time* $T$ ($T = 90$ days in the case of our data). The probability that a trader will post feedback first (second) within the time window $[0, T]$ is simply $R^k_i(T)$ ($\widetilde{R}^k_{i|j}(T|t_0)$ respectively), i.e. the relevant CDF evaluated at time

$T$. Accordingly, the probability that we will observe no feedback within $[0, T]$ is simply $1 - R_i^k(T)$ ($1 - \widetilde{R}_{i|j}^k(T|t_0)$ resp.).

Model C has 25 equations and $33 + 6P$ unknowns, where $P$ denotes the number of parameters of each failure time distribution $z(t; \cdot)$. Unknown parameters include Model A's original 15 parameters, 18 newly introduced hazard rate multipliers $\alpha_{i|j}^k$ and 6 parameter vectors $\theta_i^k$ ($i = G, M, B$, $k = b, s$) of failure time distributions $z(t; \theta_i^k)$. The size $P$ of vectors $\theta_i^k$ depends on the parametric family chosen. In contrast to Models A and B, which were finite-dimensional systems of simultaneous equations, the equations of Model C collectively define a continuous probability distribution $f(i_b i_s, t_b, t_s)$ (i.e. an an *infinite*-dimensional mathematical object) where $i_b i_s$ is the feedback pattern observed and $t_b, t_s$ are the buyer and seller feedback submission times respectively.[6] Intuitively, an infinite-dimensional model has a unique specification in terms of any finite number of scalar parameters, provided that these parameters are "sufficiently independent" from one another. This specific notion of independence is formally captured by the rank of the Fisher information matrix. Specifically, Rothenberg (1971) and Bowden (1973) have shown that the (local) identifiability of a parameter vector $\theta$ in the context of a stochastic model $f(x; \theta)$ can be established by testing that the Fisher information matrix

$$H(\theta) = [h_{ij}(\theta)] = E_x \left[ \frac{\partial log f}{\partial \theta_i} \frac{\partial log f}{\partial \theta_j} \right] \tag{7}$$

is non-singular at the parameter estimate $\hat{\theta}$. Since the Fisher matrix usually does not have a closed-form representation, in practice such a test is performed ex-post, that is, after an estimate has been derived using some statistically sound estimation method. Estimation is, in principle, feasible if the number of distinct observations is greater than the number of parameters. Since Model C takes into account the times of feedback submission in each individual auction, estimation makes use of the full set of 51,062 distinct observations in our data.

4.2. **Model estimation and results.** Each of the 51,062 auctions in our data set can be equivalently expressed as a tuple $(k, j_b, j_s, t_b, t_s)$, where $k \in \{b, s, 0\}$ denotes the trader who posts the first rating (0 indicates that no trader posts feedback), $j_b, j_s \in \{+, -, 0, S\}$ denote the type of rating posted by the buyer and seller respectively ($S$ indicates silence), and $t_b, t_s \in [0, T]$ denote the time

[6]It is straightforward (though tedious) to verify that the 25 equations of Model C sum to one and that for any feedback pattern $(i_b, i_s) \in \{+, 0, -, S\}^2$ and any pair of submission times $(t_b, t_s) \in ([0, T] \cup \oslash)^2$ (where $\oslash$ denotes no feedback) there is exactly one model equation that provides the corresponding observation probability.

of each trader's feedback relative to the closing time of the corresponding auction (a value of 0 indicates that the corresponding trader did not post feedback). Since our data set only includes ratings posted up until 90 days after the auction's closing, $T = 90$.

We model time-to-feedback densities $z(t;\theta_i^k)$ using lognormal distributions. Lognormal distributions are commonly used to model a wide range of failure time distributions (Limpert et al. 2001); for appropriate parameter ranges such distributions can approximate very well the empirical distributions of feedback posting times in our data (Figure 1). A lognormal distribution is fully defined by a parameter vector that includes two components: a location parameter $\mu \in \mathbb{R}$ and a scale parameter $\varsigma \in \mathbb{R}^+$. Since our model assumes that time-to-feedback densities are conditional on the type of trader and the type of outcome, this step introduces $2 \times 3 \times 2 = 12$ additional parameters to our model. The final model, thus, has 45 unknown parameters:

- 9 joint outcome probability parameters $\pi_{i_b i_s}$
- 6 "first mover" reporting probabilities $\rho_i^k$ (probabilities of reporting feedback conditional on the partner not having posted a rating)
- 18 hazard rate multipliers $\alpha_{i|j}^k$ indicating how the partner's feedback modifies a trader's subsequent propensity to rate
- 6 lognormal distribution location parameters $\mu_i^k$ (one parameter per trader type and outcome type)
- 6 lognormal distribution scale parameters $\varsigma_i^k$ (one parameter per trader type and outcome type)

Estimates of all parameters can be obtained by maximizing the log-likelihood function:

$$\mathcal{L} = \sum_{n=1}^{51{,}062} log\left(f_{j_b(n)j_s(n)}^{k(n)}(t_b(n), t_s(n)|T = 90)\right) \tag{8}$$

where each $f_{j_b j_s}^k(\cdot)$ is the right-hand side of the corresponding Model C equation. Parameter estimates must satisfy the constraints:

$$\pi_{i_b i_s}, \rho_i^k \in [0,1], \quad \alpha_{i|j}^k \in \mathbb{R}^+, \quad \mu_i^k \in \mathbb{R}, \quad \varsigma_i^k \in \mathbb{R}^+ \text{ and } \sum_{i_b, i_s} \pi_{i_b i_s} = 1 \tag{9}$$

Table 6 summarizes the properties of the parameter estimates that solve the above constrained maximization problem. Identifiability was established by numerically verifying that the Fisher information matrix (7) has a non-zero determinant at the parameter estimate. As before, to facilitate the interpretation of our results, Table 6 lists the estimates of the marginal probabilities of trader satisfaction $\pi_i^k$ in lieu of the joint outcome probabilities $\pi_{i_b i_s}$. We also omit the estimates of the 12 parameters $\mu_i^k$, $\varsigma_i^k$ since their interest is secondary to the purposes of this study.

In the rest of this section we discuss the most important insights provided by these results.

*Trader satisfaction and first mover reporting probabilities.* Trader satisfaction probabilities estimated by Model C are in line with those estimated by Model A. Model C estimates that, on average, buyers observe good outcomes 78.9% of the time, mediocre outcomes 20.4% of the time and bad outcomes only 0.7% of the time. The corresponding figures for sellers are 85.7%, 13.7% and 0.6% respectively. Observe that the estimates of both good and bad outcomes are a little lower than the respective estimates of Model A for both buyers and sellers. We attribute the difference in Model C's ability to more accurately estimate a trader's reporting probabilities before and after receiving a rating from his partner. Note, also, that since Model C is making use of finer-grained data, it allows us to obtain tighter confidence intervals than Model A.

In terms of first mover reporting probabilities, the most notable feature of Table 6 are the strikingly high estimates of very dissatisfied traders' propensities to report negative feedback ($\rho_B^b$, $\rho_B^s$). The corresponding 95% confidence intervals are still quite broad, so comparison with the corresponding propensities to report satisfactory outcomes is statistically ambiguous. Observe, however, that the lower bounds of these intervals (51.5% for buyers; 68.2% for sellers) allude to substantial reporting levels of "very bad" outcomes.

*Second mover reporting probabilities.* The principal new feature of Model C is that it allows us to assess how partner feedback modifies a trader's subsequent reporting behavior. Specifically, our model assumes that a partner's rating action $j$ multiplies trader $k$'s subsequent hazard rate of reporting $i$ by an unknown factor $\alpha_{i|j}^k$, which is to be estimated together with all other model parameters.

Table 6 lists the ML estimates of all 18 hazard rate multipliers $\alpha_{i|j}^k$ defined in this manner. The interpretation of these estimates is based on the following reasoning: If the 95% confidence interval of a given $\alpha_{i|j}^k$ falls entirely above (below) 1 then our model provides statistical evidence (at the 95% level) that submission of feedback $j$ by the partner increases (decreases) trader $k$'s subsequent propensity to report outcome $i$. In contrast, if the corresponding 95% confidence interval contains

| Parameter | ML Estimates | | Confidence Intervals | | |
|---|---|---|---|---|---|
| | Mean | Std. Error | 2.50% | Median | 97.50% |
| $\pi_G^b = \pi_{GG} + \pi_{GM} + \pi_{GB}$ | 0.789 | 0.003 | 0.783 | 0.789 | 0.796 |
| $\pi_M^b = \pi_{MG} + \pi_{MM} + \pi_{MB}$ | 0.204 | 0.003 | 0.198 | 0.204 | 0.210 |
| $\pi_B^b = \pi_{BG} + \pi_{BM} + \pi_{BB}$ | 0.007 | 0.001 | 0.005 | 0.006 | 0.009 |
| $\pi_G^s = \pi_{GG} + \pi_{MG} + \pi_{BG}$ | 0.856 | 0.002 | 0.851 | 0.856 | 0.861 |
| $\pi_M^s = \pi_{GM} + \pi_{MM} + \pi_{BM}$ | 0.137 | 0.003 | 0.132 | 0.137 | 0.143 |
| $\pi_B^s = \pi_{GB} + \pi_{MB} + \pi_{BB}$ | 0.006 | 0.001 | 0.005 | 0.006 | 0.009 |
| $\rho_G^b$ | 0.677 | 0.005 | 0.666 | 0.677 | 0.687 |
| $\rho_M^b$ | 0.026 | 0.003 | 0.020 | 0.026 | 0.033 |
| $\rho_B^b$ | 0.786 | 0.125 | 0.515 | 0.806 | 0.970 |
| $\rho_G^s$ | 0.819 | 0.003 | 0.812 | 0.819 | 0.827 |
| $\rho_M^s$ | 0.027 | 0.011 | 0.014 | 0.023 | 0.056 |
| $\rho_B^s$ | 0.872 | 0.075 | 0.682 | 0.880 | 0.978 |
| $\alpha_{G\|+}^b$ | **3.279** | 0.044 | 3.193 | 3.279 | 3.361 |
| $\alpha_{M\|+}^b$ | **0.366** | 0.067 | 0.253 | 0.360 | 0.514 |
| $\alpha_{B\|+}^b$ | **2.533** | 0.699 | 1.330 | 2.446 | 4.030 |
| $\alpha_{G\|0}^b$ | **0.260** | 0.130 | 0.065 | 0.240 | 0.574 |
| $\alpha_{M\|0}^b$ | 1.418 | 1.046 | 0.067 | 1.239 | 3.907 |
| $\alpha_{B\|0}^b$ | 1.177 | 0.947 | 0.039 | 0.962 | 3.379 |
| $\alpha_{G\|-}^b$ | 0.394 | 0.371 | 0.072 | 0.272 | 1.458 |
| $\alpha_{M\|-}^b$ | 1.834 | 0.873 | 0.452 | 1.709 | 3.845 |
| $\alpha_{B\|-}^b$ | 0.959 | 0.740 | 0.269 | 0.737 | 3.204 |
| $\alpha_{G\|+}^s$ | **3.648** | 0.056 | 3.537 | 3.648 | 3.763 |
| $\alpha_{M\|+}^s$ | **0.143** | 0.059 | 0.052 | 0.137 | 0.276 |
| $\alpha_{B\|+}^s$ | **0.252** | 0.228 | 0.056 | 0.176 | 0.949 |
| $\alpha_{G\|0}^s$ | **0.435** | 0.095 | 0.266 | 0.430 | 0.647 |
| $\alpha_{M\|0}^s$ | **4.035** | 1.046 | 2.145 | 4.022 | 6.179 |
| $\alpha_{B\|0}^s$ | **3.308** | 1.033 | 1.517 | 3.250 | 5.631 |
| $\alpha_{G\|-}^s$ | 0.621 | 0.413 | 0.149 | 0.512 | 1.731 |
| $\alpha_{M\|-}^s$ | 1.495 | 0.858 | 0.241 | 1.319 | 3.524 |
| $\alpha_{B\|-}^s$ | **5.276** | 0.919 | 3.622 | 5.278 | 7.272 |

TABLE 6. Maximum likelihood estimates of Model C parameters. Boldface indicates hazard rate multipliers $\alpha_{i|j}^k$ that were found to be significantly different than one.

1 then we do *not* find statistically significant evidence that submission of feedback $j$ affects the trader's subsequent reporting behavior. Table 6 lists in boldface all hazard rate multipliers $\alpha_{i|j}^k$ that were found to be *significantly different than one* according to this definition. The following paragraphs discuss what these results mean.

*Impact of seller's feedback on buyer's behavior.* Receipt of positive feedback from a seller appears to substantially increase the average buyer's propensity to report good and bad outcomes and to decrease her propensity to report mediocre outcomes. Buyers, thus, appear to return the "favor" of positive feedback by posting positive feedback for the seller with increased probability (when satisfied) and by withholding the reporting of neutral feedback (when mildly dissatisfied). On the other hand, positive feedback does not appear to be enough to appease a buyer who has experienced a bad outcome. Instead, the removal of the threat of retaliation (the seller can only post feedback once) appears to embolden dissatisfied buyers who are, then, more likely to post negative feedback for the seller.

Receipt of neutral feedback from the seller appears to decrease the average buyer's propensity to report good outcomes. Buyers who receive neutral feedback, thus, exhibit a form of negative reciprocation, withholding a positive rating that they would otherwise be likely to post for the seller. Interestingly, receipt of negative feedback does not have any statistically significant impact on the average buyer's subsequent reporting behavior.

In summary, we find strong evidence of positive reciprocation but only mild evidence of negative reciprocation on the part of eBay buyers. Buyers also appear to be sensitive to the possibility of seller retaliation, and thus, more likely to report bad outcomes *after* the seller has posted feedback.

*Impact of buyer's feedback on seller's behavior.* Receipt of positive feedback from a buyer appears to substantially increase the average seller's propensity to report good outcomes and to decrease his propensity to report mediocre and bad outcomes. Seller behavior, thus, exhibits strong positive reciprocation. In addition, sellers appear to be willing to "forgive" delinquent (i.e. late-paying or non-paying) buyers in exchange for a positive rating.

Receipt of neutral feedback from a buyer decreases the seller's propensity to report good outcomes and increases his propensity to report mediocre and bad outcomes. Seller behavior, thus, exhibits strong negative reciprocation to neutral feedback. Finally, receipt of a negative rating substantially increases the seller's propensity to report bad outcomes but does not appear to have a statistically significant impact on his propensity to report positive and mediocre outcomes.

In summary, evidence of reciprocal behavior is even stronger in the case of eBay sellers. Like buyers, eBay sellers appear to respond positively to positive feedback. Furthermore, eBay sellers appear to be more sensitive to unfavorable feedback than eBay buyers. The latter finding is intuitive: the adverse impact of negative feedback is more severe for sellers than it is for buyers. Sellers,

| Second mover | Outcome experienced by second mover | First mover's feedback *Positive* | *Neutral* | *Negative* |
|---|---|---|---|---|
| | | Impact on second mover's propensity to report | | |
| Buyer | *Good* | + | - | |
| | *Mediocre* | - | | |
| | *Bad* | + | | |
| Seller | *Good* | + | - | |
| | *Mediocre* | - | + | |
| | *Bad* | - | + | + |

Legend: + increases propensity to report - decreases propensity to report

Table 7. Summary of how the first mover's feedback affects the the second mover's subsequent propensity to report the transaction outcome she has observed.

therefore, have a higher incentive to develop a reputation of negative reciprocation that would serve to discourage a fraction of dissatisfied buyers from posting unfavorable feedback.

Table 7 summarizes the above findings, providing an easy reference of how the first mover's feedback affects the the second mover's subsequent propensity to report the transaction outcome she has observed.

## 5. Obtaining transaction-specific estimates

The models introduced in the previous two sections derive population-level estimates of *average* trader satisfaction and feedback reporting probabilities. On a market as heterogeneous as eBay, it is plausible to assume that these quantities exhibit substantial variance among traders and transaction types. Traders will, thus, often prefer to obtain *transaction-specific* estimates of the risks of trading with a specific partner.

Enhancing our modeling technology to obtain transaction-specific estimates is straightforward. The essential step is to replace our original model's unknown parameters by regression equations that describe how the relevant parameter relates to salient attributes of the trader, the partner, and the transaction of interest. The unknown parameters of the augmented model are the coefficients that relate our newly introduced covariates to the original outcome and reporting probabilities. Once estimates of these coefficients have been obtained, we can use the regression equations to obtain transaction-specific outcome and reporting probability estimates for any given covariate vector.

The rest of the section briefly illustrates this idea in the context of our eBay data set. Consistent with past empirical research in online auctions (see Bajari and Hortacsu 2004; Resnick et al. 2006 for surveys), it is plausible to assume that eBay's outcome and reporting probabilities are dependent on the following properties of the buyer, seller and transaction:

- *Trader's feedback score.* A trader's feedback score is the sum of a trader's positive minus negative ratings over her entire "career" on eBay. Given the rarity of negative ratings, a trader's feedback score should primarily be viewed as a proxy of the trader's experience on eBay as opposed to a measure of the trader's "quality." Nevertheless, it is reasonable to expect that, by virtue of having "survived" long enough, traders with high feedback scores are most likely honest and adequately competent.
- *Trader's total number of negative ratings.* A trader's total number of negative ratings is a measure of the trader's "quality." If the feedback mechanism works well, one expects trading risks to be higher when trading with partners who have received more negative ratings in the past.
- *Transaction value.* Both the probability of trader satisfaction, as well as their reporting behavior are likely to be conditioned on the value of the item being traded. High value items are typically more complex, increasing the probability that something might go wrong.
- *Number of bidders.* The number of bidders competing for an item indicates the level of demand for that item. It is plausible to expect that it may thus influence a trader's satisfaction and reporting behavior.

We now develop and fit a model that can be used to derive transaction-specific estimates of trader satisfaction and reporting probabilities as a function of the buyer and seller feedback scores (*bscore*, *sscore*), number of negative ratings received so far (*bnegs*, *snegs*), transaction value (*price*) and total number of bids from unique bidders (*bids*). Our model will be constructed by augmenting Model A. Augmentation of Model C leads to qualitatively similar results.

To apply the method we substitute each of Model A's 15 unknown parameters by the following regression equation:

$$\begin{aligned}\text{logit}(x_i) = \beta_{0i}& \\ &+\beta_{1i}\log(bscore) + \beta_{2i}\log(bnegs) + \beta_{3i}\log(sscore) \\ &+\beta_{4i}\log(snegs) + \beta_{5i}\log(price) + \beta_{6i}\log(bids)\end{aligned} \tag{10}$$

where $x_i$ represents one of our model's original parameters and $\text{logit}(x_i) = \log(x_i/(1-x_i))$. The use of the logit transformation allows coefficients $\beta_{ki}$ to remain unconstrained, while ensuring that the resulting $x_i$ always falls between 0 and 1.

| | *bscor* | *sscor* | *price* | *bids* | *bnegs* | *snegs* |
|---|---|---|---|---|---|---|
| $\pi_G^b$ | + | + | - | | - | - |
| $\pi_B^b$ | - | - | + | | + | + |
| $\pi_G^s$ | + | + | - | | - | - |
| $\pi_B^s$ | - | - | + | + | + | + |
| $\rho_G^b$ | + | - | | | - | + |
| $\rho_M^b$ | | - | | | | |
| $\rho_B^b$ | | - | | | | + |
| $\rho_G^s$ | - | + | | | + | - |
| $\rho_M^s$ | | | | - | | - |
| $\rho_B^s$ | | | | | + | |

(a) Signs of significant coefficients

| | *bscor* | *sscor* | *price* | *bids* | *bnegs* | *snegs* |
|---|---|---|---|---|---|---|
| $\pi_G^b$ | **0.0046**<br>**(0.0006)** | **0.0016**<br>**(0.0002)** | **-0.0004**<br>**(0.0001)** | -0.0002<br>(0.0002) | **-0.1995**<br>**(0.0113)** | **-0.1332**<br>**(0.0053)** |
| $\pi_B^b$ | **-0.0094**<br>**(0.0009)** | **-0.0021**<br>**(0.0003)** | **0.0004**<br>**(0.0001)** | -0.0001<br>(0.0059) | **0.2569**<br>**(0.0142)** | **0.1524**<br>**(0.0065)** |
| $\pi_G^s$ | **0.0096**<br>**(0.0006)** | **0.0006**<br>**(0.0001)** | **-0.0002**<br>**(0.0001)** | -0.0060<br>(0.0038) | **-0.1699**<br>**(0.0086)** | **-0.1103**<br>**(0.0059)** |
| $\pi_B^s$ | **-0.0226**<br>**(0.0023)** | **-0.0030**<br>**(0.0010)** | **0.0003**<br>**(0.0001)** | **0.0287**<br>**(0.082)** | **0.0563**<br>**(0.0134)** | **0.1573**<br>**(0.0090)** |
| $\rho_G^b$ | **0.0051**<br>**(0.0005)** | **-0.0005**<br>**(0.0001)** | 0.0000<br>(0.0001) | -0.0036<br>(0.0030) | **-0.1238**<br>**(0.0108)** | **0.0489**<br>**(0.0089)** |
| $\rho_M^b$ | 0.0021<br>(0.0011) | **-0.0013**<br>**(0.0007)** | -0.0006<br>(0.0005) | 0.0277<br>(0.0147) | -0.0270<br>(0.0321) | 0.0212<br>(0.0134) |
| $\rho_B^b$ | 0.0058<br>(0.0043) | **-0.0096**<br>**(0.0027)** | 0.0004<br>(0.0020) | 0.0377<br>(0.0468) | -0.0606<br>(0.1051) | **3.1440**<br>**(0.6160)** |
| $\rho_G^s$ | **-0.0013**<br>**(0.0002)** | **0.0018**<br>**(0.0001)** | 0.0001<br>(0.0001) | 0.0000<br>(0.0032) | **0.0203**<br>**(0.0104)** | **-0.1476**<br>**(0.0054)** |
| $\rho_M^s$ | 0.0036<br>(0.0027) | -0.0003<br>(0.0009) | -0.0008<br>(0.0009) | **-0.1000**<br>**(0.0357)** | -0.0373<br>(0.0500) | **-0.1320**<br>**(0.0603)** |
| $\rho_B^s$ | -0.0030<br>(0.0041) | 0.0005<br>(0.0017) | 0.0004<br>(0.0015) | -0.0382<br>(0.0542) | **5.3240**<br>**(0.5683)** | -0.0611<br>(0.0645) |

(b) Coefficient magnitudes followed by standard errors in parentheses
(boldface indicates significance)

TABLE 8. Coefficient estimates of transaction-specific model.

Substituting each of Model A's 15 original parameters by the corresponding equation (10) we obtain an augmented model that has $15 \times 7 = 105$ unknown coefficients $\beta_{ki}$. Coefficient estimates can be estimated, as in all previous cases, using the maximum likelihood method.

Table 8 lists the sign and magnitudes of all coefficients $\beta_{ki}$, highlighting those that were found to be statistically significant (at the 95% level).[7] As usual, the table lists the coefficients that predict a trader's marginal probabilities of satisfaction $\pi_G^k$ and dissatisfaction $\pi_B^k$ in lieu of the, less intuitive,

[7]Statistical significance in this context means that the 95% confidence interval of a coefficient's ML estimate does not contain zero.

joint outcome probabilities $\pi_{i_b i_s}$. The following paragraphs briefly discuss some insights provided by these results. Our explanations are meant to be tentative but nevertheless reveal some interesting and intuitive relationships.

*Trader satisfaction.* Our results indicate that a buyer's marginal probability of satisfaction ($\pi_G^b$) exhibits positive correlation with the seller's feedback score *sscore* and negative correlation with the seller's number of negative ratings *snegs*. These results suggest that, despite the presence of reporting bias, the reputation metrics published by eBay's simple feedback mechanism are statistically significant in predicting buyer satisfaction. Interestingly, the absolute magnitude of the coefficient of *snegs* is substantially higher than that of the coefficient of *sscor*, indicating that the seller's number of negative ratings is a more important predictor of buyer satisfaction than the seller's eBay feedback score.

Buyer satisfaction is negatively correlated with the item's *price*; this is intuitive, since more expensive items also tend to be more complex, increasing the probability that the buyer's expectations will not be matched by what she receives.

We were somewhat intrigued to find that the buyer's *own* reputation metrics ($bscor, bnegs$) were significant in predicting her probability of satisfaction. This is consistent with our earlier remark that dissatisfaction is sometimes due to the buyer's own lack of experience with how eBay works: more experienced buyers are, thus, better able to understand what the seller promises to deliver as well as to communicate more effectively with him. Furthermore, more experienced eBay buyers are perhaps better able to "see through" aspects of a transaction that are not captured by any of our other covariates and to avoid transactions that contain subtle cues of risk.

The determinants of the probability of buyer *dis*satisfaction ($\pi_B^b$) are identical to the determinants of buyer satisfaction. Exactly the same coefficients ended up being significant, with opposite signs and similar magnitudes.

The determinants of the probability of seller satisfaction $\pi_G^s$ (dissatisfaction $\pi_B^s$) are also very similar to those of buyer satisfaction: Seller satisfaction (dissatisfaction) appears to be positively (negatively) correlated with the buyer's feedback score *bscore*, negatively (positively) correlated with the buyer's number of negative ratings *bnegs*, negatively (positively) correlated with the item's *price* and also affected by the seller's own feedback score and number of negative ratings.

We attribute the negative correlation between seller satisfaction $\pi_G^s$ and *price* to the higher probability of complications (no-paying or late-paying buyers, buyers making additional demands, etc.) in high value item transactions. Also, the higher the price, the higher the probability of *winner's curse* phenomena: situations where buyers feel that they have bid above their valuation, and as a result, might decide to back out of their commitment to buy. The winner's curse can also provide an explanation for the positive correlation between seller *dis*satisfaction $\pi_B^s$ and the number of unique bidders (*bids*) participating in an auction: auctions where many bidders compete often "push" bidders to bid above their valuation. This may, in turn, increase the probability that the winning bidder backs out of his offer to buy the item, which, as we discuss in Section 2, represents the primary cause of dissatisfied sellers.

*Reporting probabilities.* A buyer's probability of reporting good outcomes ($\rho_G^b$) is positively correlated with the buyer's own feedback score (*bscor*) and negatively correlated with the buyer's number of negative ratings (*bnegs*). A possible explanation of these relationships is that traders with more experience and a better record are also "better citizens" in terms of their participation to the feedback mechanism. A complementary explanation is that more experienced sellers have probably figured out that posting positive feedback for a buyer increases the probability that the buyer will return the favor (see Table 7) and are being, therefore, strategically proactive in posting such feedback.

Interestingly, all three buyer reporting probabilities ($\rho_G^b$, $\rho_M^b$, $\rho_B^b$) have a negative relationship with the seller's feedback score (*sscor*). This means that buyers appear to be less willing to provide feedback for experienced sellers (or, conversely, more willing to provide feedback for less experienced sellers). One tentative explanation of this empirical relationship is that it signals the presence of altruistic motives: Inexperienced sellers need positive feedback more than experienced sellers. Likewise, the rest of the community benefits more if a bad seller is "exposed" sooner rather than later. Both arguments imply that, if altruism is part of a buyer's motivation to post feedback, this motivation will be stronger if the seller has less experience.

Last, but not least, buyer reporting probabilities $\rho_G^b$, $\rho_B^b$ are positively correlated with the seller's number of negatives (*snegs*). The magnitude of the corresponding coefficient is particularly large in the case of bad outcomes, suggesting that dissatisfied buyers are substantially more willing to post negative feedback for sellers who have already received negative feedback in the past. This interesting empirical fact is consistent with Khopkar et al. (2005) "stoning" theory of eBay buyers, and represents a very different method for reaching the same conclusion.

Similar to buyers, sellers with more experience (higher *sscor*) and a better record (lower *snegs*) are more likely to report good outcomes. Also similar to buyers, dissatisfied sellers appear to be substantially more likely to post negative feedback when the relevant buyer has already received negative feedback in the past.

Somewhat curiously, the number of unique bidders (*bids*) has a statistically significant negative relationship with a seller's propensity to report a mediocre outcome. We offer a purely tentative interpetation of this last finding: assuming that the source of most mediocre outcomes for a seller is a buyer backing out of a transaction, the presence of several other interested bidders might make it easier for the seller to find an alternative seller for the item, thus lessening his dissatisfaction.

## 6. Managerial implications and Conclusions

Internet-enabled feedback mechanisms offer society a tremendous potential to reduce information asymmetries and, thus, to increase the efficiency of electronic and traditional markets (Bakos and Dellarocas 2005). The value of feedback mechanisms is, of course, only as good as the quality of information that is reported to them by their participants. In most practical settings, transaction outcomes are privately observed and voluntarily reported. Voluntary reporting, in turn, introduces the potential for reporting bias. Reporting bias arises when a person's propensity to report a privately observed outcome to a public feedback mechanism is conditioned on the type of outcome: some types of outcomes then get reported more often than others, distorting the distribution of public feedback relative to the distribution of the private transaction outcomes and potentially leading the users of this feedback to erroneous conclusions.

Reporting bias has many and complex causes. Some of these causes have their roots on the fundamentals of human behavior: it is widely accepted that people are more willing to disclose extreme experiences than average experiences (Anderson 1998) and reluctant to transmit bad news (Tesser and Rosen 1975). Other causes might be due to fear of litigation or other forms of retaliation from a transaction partner and to distortions introduced by the feedback mechanism itself. For example, in mechanisms that reward raters on the basis of how many users have found their reviews useful, strategic reviewers might be biased towards reviewing more popular products, for which the audience (and, thus, the likely volume of "usefulness votes") is larger.

Faced with the adverse consequences of reporting bias, feedback mechanism designers and users have two non-rival paths of action. The first path is to explore to what extent changes in the design of a feedback mechanism can reduce reporting bias. There are several interesting possibilities in this direction; we will briefly discuss some of them later on in this section. The second path (and the focus of this paper) is to develop methodologies that can help mechanism users make better inferences from feedback provided by today's imperfect mechanisms. The practical value of this alternative path is reinforced by the belief that, although there will always be room for mechanism design improvements, there will, likewise, always be cases where reporting bias cannot be fully eliminated or where incentive conflicts among stakeholders will make it politically impossible to implement improved mechanisms.[8]

Consistent with the above perspective, the primary contribution of this paper is a methodology that allows users of bidirectional feedback mechanisms to derive unbiased estimates of the distribution of privately observed transaction outcomes from a sample of, potentially biased, public feedback. The method also provides estimates of the mechanism participants' reporting behavior and can, thus, reveal the extent to which reporting bias is present in a given setting. Our approach is based on extracting information from the temporal order of buyer and seller feedback submission as well as from the relative incidence of transactions where one or both traders choose to remain silent.

We apply our method to a large data set of eBay feedback and obtain quantitative estimates of eBay transaction outcomes and eBay traders' reporting behavior. To the best of our knowledge we are the first to provide concrete numerical estimates of the degree to which several phenomena that have been hypothesized in the literature are present on eBay. Specifically, we confirm the widespread belief that eBay traders are more likely to post feedback when satisfied than when mildy unsatisfied. (Reporting probabilities go up again when traders are very dissatisfied.) Furthermore, we provide statistically rigorous evidence for the presence of positive and negative reciprocation among eBay traders. Last but not least, our results allow us to "see through" the, overwhelmingly positive, eBay feedback and derive what, we believe, are more realistic estimates of the risks associated with eBay transactions.

Our method can be augmented with covariates to derive more accurate, transaction-specific estimates of transaction risks. We illustrate how this can be done in the context of our eBay data set, deriving regression equations that can predict the probability of an eBay trader's satisfaction on a

[8]For example, eBay has found that many proposed changes to its feedback mechanism were met with fierce resistance from its "power sellers" who use their eBay feedback score as a marketing tool.

specific transaction, given his feedback profile information, his partner's feedback profile information, the number of unique bidders and the transaction value.

The application of our method to eBay feedback illustrates both its power and the wealth of insights that it is capable of producing. Our method is quite general and can be applied, with a little fine-tuning, to many feedback mechanisms (not necessarily online) where both parties of a transaction have the right to report their satisfaction. We, therefore, believe that the method can have important practical implications in several industries where the application of online feedback mechanisms is currently hindered by reporting bias considerations. For example, several current attempts to develop online feedback mechanisms for physicians are hampered by the threat of physician lawsuits against unfavorable comments posted by patients (Kesmodel 2005). If such lawsuits prove successful, it is plausible to expect that the fear of litigation will reduce the propensity of dissatisfied patients to report their physician online and, therefore, that, just as eBay, such sites will end up having deceivingly high percentages of positive feedback. Our method can be applied to such settings to extract information from a patient's choice to remain silent and to derive unbiased estimates of the rate of satisfaction associated with specific physicians.[9]

The results of our work provide useful insights that can inform the design of future feedback mechanisms. We show that reciprocity (both positive and negative) is a powerful driver of people's feedback reporting behavior. The fear of retaliation has been widely publicized as an important reason behind people's reluctance to report negative outcomes on eBay (Klein et al. 2005). Motivated by such remarks, some authors have proposed that eBay should allow only the buyer to rate the seller (Chwelos and Dhar 2006) or that it should simultaneously reveal both partners' ratings (Reichling 2004). Our results suggest that the situation is potentially more complex. Although we find evidence of negative reciprocation, our results also establish powerful evidence of *positive* reciprocation. Whereas the fear of negative reciprocation might discourage reporting of bad outcomes, positive reciprocation *increases* positive feedback reporting levels, which, in turn, increases market efficiency (Dellarocas 2005). It is, in fact, plausible to argue that the expectation of positive reciprocation from one's partner might be an important reason why eBay traders report positive transaction outcomes with such high probability (around 70% in our data), despite early predictions that online feedback would be underprovided (Avery, Resnick and Zeckhauser 1999). Changing the

[9]For our method to apply, the reputation mechanism must be able to obtain records of all patient-physician transactions, whether rated or not. It is plausible to imagine that reputation mechanisms operated by insurance companies would have access to this information.

mechanism's design to either allow only one partner to rate the other or to simultaneously reveal both partners' ratings would indeed remove the fear of retaliation; it would, however, also remove the participation incentives that are related to positive reciprocation. Unless we experimentally test how people respond to such mechanism changes, it is difficult to tell whether their overall impact would be to increase or to decrease a mechanism's effectiveness.

The one unequivocal message that our study delivers to feedback mechanism designers is that a trader's decision to *not* post online feedback carries important information that can assist the users of feedback to make more reliable inferences. The majority of today's feedback mechanisms does not publicly disclose the number of *silent transactions* (i.e., transactions for which no feedback was posted by one or both partners). We argue that such information should become a part of a trader's feedback profile on eBay and other feedback mechanisms.

We conclude the paper by pointing out some caveats and limitations of our method, together with associated opportunities for future research. First, due to identifiability constraints, our models work in settings where traders exhibit reporting bias but not strategic misreporting. The prevalence of "one-time" trading relationships on eBay makes extensive strategic misreporting improbable but certainly not impossible. The empirical results obtained in the context of eBay should, therefore, be interpreted with caution. Second, the maximum likelihood estimation (MLE) method used to estimate our models is appropriate in large-sample settings but known to generate biased estimates when samples are small. We do not expect this to be an issue in online settings, where data sets tend to be large. On the other hand, the estimation of our models in small sample settings might benefit from the use of specialized small-sample MLE extensions (see, for example, Sprott 1980). Finally, our work estimates the presence of reporting bias but does not attempt to quantify the impact of such bias on market efficiency. An important next step would be to study the extent to which traders are aware or unaware of reporting bias in a given context and take it into account in their bidding decisions. Such work will help assess the social cost of reporting bias and, thus, the benefit of developing better feedback mechanisms that are capable of compensating for it.

## References

Anderson, E. W. (1998) Customer satisfaction and word of mouth. *Journal of Service Research* 1 (1) 5-17.

Avery, C., Resnick, P., R. Zeckhauser (1999) The Market for Evaluations. *American Economics Review* 89 (3) 564-584.

Bajari, P. and Hortacsu, A. (2004) Economic Insights from Internet Auctions. *Journal of Economic Literature* 42 (2) 457-486.

Bakos, Y. and Dellarocas, C. (2005) Online Reputation and Litigation as Mechanisms for Quality Assurance: Strategic Implications for Profits and Efficiency. Working Paper.

Bollen, K.A. (1989) *Structural equations with latent variables.* New York: Wiley.

Bowden, R. (1973) The theory of parametric identification. *Econometrica* 41 (6) 1069-1074.

Chwelos, P. and Dhar, T. (2006) Caveat Emptor: Differences in Online Reputation Mechanisms. Working Paper, Sauder School of Business, University of British Columbia.

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

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

Elmore, R. T. and Wang, S. (2003) Identifiability and estimation in multinomial mixture models. Technical Report #03-04, Department of Statistics, Penn State University.

Fehr, E. and Gachter, S. (2000) Fairness and Retaliation: The Economics of Reciprocity. *Journal of Economic Perspectives* 14 (3) 159-181.

Greene, W. H. (2002) Econometric Analysis (5th Edition). Prentice-Hall, Upper Saddle River, NJ.

Hu, N., Pavlou P. A. and Zhang J. (2006) Can Online Reviews Reveal a Products True Quality? Empirical Findings and Analytical Modeling of Online Word-of-Mouth Communication. *Proc. ACM EC 06 Conference on Electronic Commerce*, Ann Arbor, Michigan.

Kauffman, R. J., C. A. Wood (2005) Doing Their Bidding: An Empirical Examination of Factors that Affect a Buyer's Utility in Internet Auctions. *Information Technology and Management* 7 (3) 171-190.

Kesmodel, D. (2005) As Angry Patients Vent Online, Doctors Sue to Silence Them. *The Wall Street Journal*, September 14, 2005.

Khopkar, T., Li, X., and Resnick, P. (2005) Self-Selection, Slipping, Salvaging, Slacking, and Stoning: the Impacts of Negative Feedback at eBay. *Proc. ACM EC 05 Conference on Electronic Commerce*, Vancouver, Canada.

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

Limpert, E., Stahel, W. and M. Abbt (2001) Log-normal Distributions across the Sciences: Keys and Clues. *BioScience* 51 (5) 341-352.

Reichling, F. (2004) Effects of Reputation Mechanisms on Fraud Prevention in eBay Auction. Working Paper. Stanford University.

Resnick, P., R. Zeckhauser (2002) Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay's Reputation System. Michael R. Baye, ed. *The Economics of the Internet and E-Commerce (Advances in Applied Microeconomics, Vol. 11)*, JAI Press.

Resnick, P., Zeckhauser, R., Swanson, J., K. Lockwood (2006) The Value of Reputation on eBay: A Controlled Experiment. *Experimental Economics* 9 (2) 79-101.

Rothenberg, T. J. (1971) Identification in parametric models. *Econometrica* 39 (3) 577-591.

Sprott, D. A. (1980) Maximum Likelihood in Small Samples: Estimation in the Presence of Nuisance Parameters. *Biometrika* 67 (3) 515-523.

Teicher, H. (1967) Identifiability of mixtures of product measures. *Ann. Math. Statist.* 38, 1300-1302.

Tesser, A. and Rosen S. (1975) The Reluctance to Transmit Bad News, in *Advances In Experimental Social Psychology*, L.Berkowitz ed. New York: Academic Press, 193-232.

## Proofs

**Proposition 1.** The stochastic process that leads trader $k$ to post feedback $j \in \{1, .., M+1\}$ conditional on having observed outcome $i$ is equivalent to a single trial of a multinomial distribution whose $M+1$ distinct outcomes have probabilities $\rho^k_{j|i}$. Model (1) can, thus, be equivalently expressed as a $N^2$-component finite mixture of the product of two single-trial multinomial distributions where the $N^2$ mixing probabilities correspond to $\pi_{i_b i_s}$, the first multinomial has $M+1$ outcomes with respective probabilities $\rho^b_{j_b|i_b}$ and the second multinomial also has $M+1$ outcomes with respective probabilities $\rho^s_{j_s|i_s}$. Teicher (1967) has shown that finite mixtures of products of a given distribution are identifiable if and only if finite mixtures of the base distribution are identifiable. However, Elmore and Wang (2003) have shown that finite mixtures of multinomial distributions are identifiable if and only if the number of mixture components $\mu$ and the number of multinomial trials $\tau$ satisfy $\tau \geq 2\mu - 1$. In our context, $\tau = 1$, $\mu = N^2$ and the identifiability condition fails for all $N \geq 2$. (Note that the condition is independent of the number of each trial's distinct outcomes $M+1$.) Therefore, model (1) is not identifiable.

**Proposition 2.** Assumption (A1) implies that $M = N$. If, following each of the $N$ possible outcomes, traders either report the the truth of stay silent, this implies that there are $N$ unknown reporting probabilities for each of the two traders. Together with the $N^2 - 1$ independent outcome probabilities, the total number of model unknowns is, thus, $N^2 - 1 + 2N$. If $M = N$ the number of independent model equations is $(N+1)^2 - 1 = N^2 + 2N$. Since $N^2 + 2N > N^2 + 2N - 1$ the order condition is satisfied for all $N \geq 2$. Proving the validity of the rank condition is tedious but straightforward using symbolic algebra software. Appendix IV provides MAPLE code that performs the proof.

**Proposition 3.** It is well-known that any CDF $F(t)$ relates to its hazard rate $h(t) = f(t)/(1-F(t))$ through the following relationship:

$$1 - F(t) = e^{-\int_0^t h(\tau) d\tau} \tag{11}$$

Let $\widetilde{F}(t)$ be the CDF whose hazard rate $\widetilde{h}(t)$ satisfies:

$$\widetilde{h}(t) = \begin{cases} h(t) & t < t_0 \\ \alpha h(t) & t \geq t_0 \end{cases}$$

For $t \geq t_0$ substitution into (11) gives:

$$(12) \qquad 1-\widetilde{F}(t)=e^{-\int_0^t \widetilde{h}(\tau)d\tau}=e^{-\int_0^{t_0} h(\tau)d\tau-\alpha\int_{t_0}^t h(\tau)d\tau}=e^{-\int_0^{t_0} h(\tau)d\tau}\left(\frac{e^{-\int_0^t h(\tau)d\tau}}{e^{-\int_0^{t_0} h(\tau)d\tau}}\right)^\alpha$$

Substituting the left hand side of (11) to the right hand side of (12) we obtain:

$$1-\widetilde{F}(t)=(1-F(t_0))\left(\frac{1-F(t)}{1-F(t_0)}\right)^\alpha \Rightarrow \widetilde{F}(t)=1-(1-F(t_0))\left(\frac{1-F(t)}{1-F(t_0)}\right)^\alpha,\ t\geq t_0$$

Differentiation of $\widetilde{F}(t)$ with respect to $t$ gives the form of the associated pdf for $t \geq t_0$.

## Appendix I: Model A

### Model Equations

$$
\begin{aligned}
F_{++} &= \pi_{GG}\rho_G^b\rho_G^s \\
F_{+0} &= \pi_{GM}\rho_G^b\rho_M^s \\
F_{+-} &= \pi_{GB}\rho_G^b\rho_B^s \\
\\
F_{0+} &= \pi_{MG}\rho_M^b\rho_G^s \\
F_{00} &= \pi_{MM}\rho_M^b\rho_M^s \\
F_{0-} &= \pi_{MB}\rho_M^b\rho_B^s \\
\\
F_{-+} &= \pi_{BG}\rho_B^b\rho_G^s \\
F_{-0} &= \pi_{BM}\rho_B^b\rho_M^s \\
F_{--} &= \pi_{BB}\rho_B^b\rho_B^s \\
\\
F_{+S} &= \pi_{GG}\rho_G^b(1-\rho_G^s) + \pi_{GM}\rho_G^b(1-\rho_M^s) + \pi_{GB}\rho_G^b(1-\rho_B^s) \\
F_{0S} &= \pi_{MG}\rho_M^b(1-\rho_G^s) + \pi_{MM}\rho_M^b(1-\rho_M^s) + \pi_{MB}\rho_M^b(1-\rho_B^s) \\
F_{-S} &= \pi_{BG}\rho_B^b(1-\rho_G^s) + \pi_{BM}\rho_B^b(1-\rho_M^s) + \pi_{BB}\rho_B^b(1-\rho_B^s) \\
\\
F_{S+} &= \pi_{GG}(1-\rho_G^b)\rho_G^s + \pi_{MG}(1-\rho_M^b)\rho_G^s + \pi_{BG}(1-\rho_B^b)\rho_G^s \\
F_{S0} &= \pi_{GM}(1-\rho_G^b)\rho_M^s + \pi_{MM}(1-\rho_M^b)\rho_M^s + \pi_{BM}(1-\rho_B^b)\rho_M^s \\
F_{S-} &= \pi_{GB}(1-\rho_G^b)\rho_B^s + \pi_{MB}(1-\rho_M^b)\rho_B^s + \pi_{BB}(1-\rho_B^b)\rho_B^s \\
\\
F_{SS} &= \pi_{GG}(1-\rho_G^b)(1-\rho_G^s) + \pi_{GM}(1-\rho_G^b)(1-\rho_M^s) + \pi_{GB}(1-\rho_G^b)(1-\rho_B^s) \\
&+ \pi_{MG}(1-\rho_M^b)(1-\rho_G^s) + \pi_{MM}(1-\rho_M^b)(1-\rho_M^s) + \pi_{MB}(1-\rho_M^b)(1-\rho_B^s) \\
&+ \pi_{BG}(1-\rho_B^b)(1-\rho_G^s) + \pi_{BM}(1-\rho_B^b)(1-\rho_M^s) + \pi_{BB}(1-\rho_B^b)(1-\rho_B^s)
\end{aligned}
$$

### Unknown Parameters

| # | Symbol | Description |
|---|---|---|
| 1 | $\pi_{GG}$ | Probability of good outcome for both traders |
| 2 | $\pi_{GM}$ | Probability of good outcome for buyer, mediocre for seller |
| 3 | $\pi_{GB}$ | Probability of good outcome for buyer, bad for seller |
| 4 | $\pi_{MG}$ | Probability of mediocre outcome for buyer, good for seller |
| 5 | $\pi_{MM}$ | Probability of mediocre outcome for both traders |
| 6 | $\pi_{MB}$ | Probability of mediocre outcome for buyer, bad for seller |
| 7 | $\pi_{BG}$ | Probability of bad outcome for buyer, good for seller |
| 8 | $\pi_{BM}$ | Probability of bad outcome for buyer, mediocre for seller |
| 9 | $\pi_{BB}$ | Probability of bad outcome for both traders |
| 10 | $\rho^b_G$ | Probability that a satisfied buyer will report positive feedback |
| 11 | $\rho^b_M$ | Probability that a mildly dissatisfied buyer will report neutral feedback |
| 12 | $\rho^b_B$ | Probability that a very dissatisfied buyer will report negative feedback |
| 13 | $\rho^s_G$ | Probability that a satisfied seller will report positive feedback |
| 14 | $\rho^s_M$ | Probability that a mildly dissatisfied seller will report neutral feedback |
| 15 | $\rho^s_B$ | Probability that a very dissatisfied seller will report negative feedback |

## Appendix II: Model B

### Model Equations

$$
\begin{aligned}
F^b_{++} &= \pi_{GG}\rho^b_G\widetilde{\rho}^s_{G|+}\\
F^b_{+0} &= \pi_{GM}\rho^b_G\widetilde{\rho}^s_{M|+}\\
F^b_{+-} &= \pi_{GB}\rho^b_G\widetilde{\rho}^s_{B|+}\\
F^s_{++} &= \pi_{GG}\widetilde{\rho}^b_{G|+}\rho^s_G\\
F^s_{+0} &= \pi_{GM}\widetilde{\rho}^b_{G|0}\rho^s_M\\
F^s_{+-} &= \pi_{GB}\widetilde{\rho}^b_{G|-}\rho^s_B
\end{aligned}
$$

$$
\begin{aligned}
F^b_{0+} &= \pi_{MG}\rho^b_M\widetilde{\rho}^s_{G|0}\\
F^b_{00} &= \pi_{MM}\rho^b_M\widetilde{\rho}^s_{M|0}\\
F^b_{0-} &= \pi_{MB}\rho^b_M\widetilde{\rho}^s_{B|0}\\
F^s_{0+} &= \pi_{MG}\widetilde{\rho}^b_{M|+}\rho^s_G\\
F^s_{00} &= \pi_{MM}\widetilde{\rho}^b_{M|0}\rho^s_M\\
F^s_{0-} &= \pi_{MB}\widetilde{\rho}^b_{M|-}\rho^s_B
\end{aligned}
$$

$$
\begin{aligned}
F^b_{-+} &= \pi_{BG}\rho^b_B\widetilde{\rho}^s_{G|-}\\
F^b_{-0} &= \pi_{BM}\rho^b_B\widetilde{\rho}^s_{M|-}\\
F^b_{--} &= \pi_{BB}\rho^b_B\widetilde{\rho}^s_{B|-}\\
F^s_{-+} &= \pi_{BG}\widetilde{\rho}^b_{B|+}\rho^s_G\\
F^s_{-0} &= \pi_{BM}\widetilde{\rho}^b_{B|0}\rho^s_M\\
F^s_{--} &= \pi_{BB}\widetilde{\rho}^b_{B|-}\rho^s_B
\end{aligned}
$$

$$
\begin{aligned}
F^b_{+S} &= \pi_{GG}\rho^b_G(1-\widetilde{\rho}^s_{G|+}) + \pi_{GM}\rho^b_G(1-\widetilde{\rho}^s_{M|+}) + \pi_{GB}\rho^b_G(1-\widetilde{\rho}^s_{B|+})\\
F^b_{0S} &= \pi_{MG}\rho^b_M(1-\widetilde{\rho}^s_{G|0}) + \pi_{MM}\rho^b_M(1-\widetilde{\rho}^s_{M|0}) + \pi_{MB}\rho^b_M(1-\widetilde{\rho}^s_{B|0})\\
F^b_{-S} &= \pi_{BG}\rho^b_B(1-\widetilde{\rho}^s_{G|-}) + \pi_{BM}\rho^b_B(1-\widetilde{\rho}^s_{M|-}) + \pi_{BB}\rho^b_B(1-\widetilde{\rho}^s_{B|-})\\
F^s_{S+} &= \pi_{GG}(1-\widetilde{\rho}^b_{G|+})\rho^s_G + \pi_{MG}(1-\widetilde{\rho}^b_{M|+})\rho^s_G + \pi_{BG}(1-\widetilde{\rho}^b_{B|+})\rho^s_G\\
F^s_{S0} &= \pi_{GM}(1-\widetilde{\rho}^b_{G|0})\rho^s_M + \pi_{MM}(1-\widetilde{\rho}^b_{M|0})\rho^s_M + \pi_{BM}(1-\widetilde{\rho}^b_{B|0})\rho^s_M\\
F^s_{S-} &= \pi_{GB}(1-\widetilde{\rho}^b_{G|-})\rho^s_B + \pi_{MB}(1-\widetilde{\rho}^b_{M|-})\rho^s_B + \pi_{BB}(1-\widetilde{\rho}^b_{B|-})\rho^s_B
\end{aligned}
$$

$$
\begin{aligned}
F_{SS} &= \pi_{GG}(1-\rho^b_G)(1-\rho^s_G) + \pi_{GM}(1-\rho^b_G)(1-\rho^s_M) + \pi_{GB}(1-\rho^b_G)(1-\rho^s_B)\\
&+ \pi_{MG}(1-\rho^b_M)(1-\rho^s_G) + \pi_{MM}(1-\rho^b_M)(1-\rho^s_M) + \pi_{MB}(1-\rho^b_M)(1-\rho^s_B)\\
&+ \pi_{BG}(1-\rho^b_B)(1-\rho^s_G) + \pi_{BM}(1-\rho^b_B)(1-\rho^s_M) + \pi_{BB}(1-\rho^b_B)(1-\rho^s_B)
\end{aligned}
$$

Outcome Probabilities

| | |
|---|---|
| $\pi_{GG}$ | Probability of good outcome for both traders |
| $\pi_{GM}$ | Probability of good outcome for buyer, mediocre for seller |
| $\pi_{GB}$ | Probability of good outcome for buyer, bad for seller |
| $\pi_{MG}$ | Probability of mediocre outcome for buyer, good for seller |
| $\pi_{MM}$ | Probability of mediocre outcome for both traders |
| $\pi_{MB}$ | Probability of mediocre outcome for buyer, bad for seller |
| $\pi_{BG}$ | Probability of bad outcome for buyer, good for seller |
| $\pi_{BM}$ | Probability of bad outcome for buyer, mediocre for seller |
| $\pi_{BB}$ | Probability of bad outcome for both traders |

Reporting Probabilities When Trader Rates First

| | |
|---|---|
| $\rho^b_G$ | Prob. that satisfied buyers report positive feedback if seller has not yet rated |
| $\rho^b_M$ | Prob. that mildly dissatisfied buyers report neutral feedback if seller has not yet rated |
| $\rho^b_B$ | Prob. that very dissatisfied buyers report negative feedback if seller has not yet rated |
| $\rho^s_G$ | Prob. that satisfied sellers report positive feedback if buyer has not yet rated |
| $\rho^s_M$ | Prob. that mildly dissatisfied sellers report neutral feedback if buyer has not yet rated |
| $\rho^s_B$ | Prob. that very dissatisfied sellers report negative feedback if buyer has not yet rated |

Reporting Probabilities When Trader Rates Second

(Conditional on rating received by partner)

| | | | |
|---|---|---|---|
| $\widetilde{\rho}^b_{G\|+}$ | $\widetilde{\rho}^b_{G\|0}$ | $\widetilde{\rho}^b_{G\|-}$ | Prob. that satisfied buyers report positive feedback following partner's rating |
| $\widetilde{\rho}^b_{M\|+}$ | $\widetilde{\rho}^b_{M\|0}$ | $\widetilde{\rho}^b_{M\|-}$ | Prob. that mildly dissatisfied buyers report neutral feedback following partner's rating |
| $\widetilde{\rho}^b_{B\|+}$ | $\widetilde{\rho}^b_{B\|0}$ | $\widetilde{\rho}^b_{B\|-}$ | Prob. that very dissatisfied buyers report negative feedback following partner's rating |
| $\widetilde{\rho}^s_{G\|+}$ | $\widetilde{\rho}^s_{G\|0}$ | $\widetilde{\rho}^s_{G\|-}$ | Prob. that satisfied sellers report positive feedback following partner's rating |
| $\widetilde{\rho}^s_{M\|+}$ | $\widetilde{\rho}^s_{M\|0}$ | $\widetilde{\rho}^s_{M\|-}$ | Prob. that mildly dissatisfied sellers report neutral feedback following partner's rating |
| $\widetilde{\rho}^s_{B\|+}$ | $\widetilde{\rho}^s_{B\|0}$ | $\widetilde{\rho}^s_{B\|-}$ | Prob. that very dissatisfied sellers report negative feedback following partner's rating |

## Appendix III: Model C

### Model Equations

$$
\begin{aligned}
f^b_{++}(t_b,t_s) &= \pi_{GG} r^b_G(t_b)\widetilde{r}^s_{G|+}(t_s|t_b)\\
f^b_{+0}(t_b,t_s) &= \pi_{GM} r^b_G(t_b)\widetilde{r}^s_{M|+}(t_s|t_b)\\
f^b_{+-}(t_b,t_s) &= \pi_{GB} r^b_G(t_b)\widetilde{r}^s_{B|+}(t_s|t_b)\\
f^s_{++}(t_b,t_s) &= \pi_{GG} \widetilde{r}^b_{G|+}(t_b|t_s) r^s_G(t_s)\\
f^s_{+0}(t_b,t_s) &= \pi_{GM} \widetilde{r}^b_{G|0}(t_b|t_s) r^s_M(t_s)\\
f^s_{+-}(t_b,t_s) &= \pi_{GB} \widetilde{r}^b_{G|-}(t_b|t_s) r^s_B(t_s)
\end{aligned}
$$

$$
\begin{aligned}
f^b_{0+}(t_b,t_s) &= \pi_{MG} r^b_M(t_b)\widetilde{r}^s_{G|0}(t_s|t_b)\\
f^b_{00}(t_b,t_s) &= \pi_{MM} r^b_M(t_b)\widetilde{r}^s_{M|0}(t_s|t_b)\\
f^b_{0-}(t_b,t_s) &= \pi_{MB} r^b_M(t_b)\widetilde{r}^s_{B|0}(t_s|t_b)\\
f^s_{0+}(t_b,t_s) &= \pi_{MG} \widetilde{r}^b_{M|+}(t_b|t_s) r^s_G(t_s)\\
f^s_{00}(t_b,t_s) &= \pi_{MM} \widetilde{r}^b_{M|0}(t_b|t_s) r^s_M(t_s)\\
f^s_{0-}(t_b,t_s) &= \pi_{MB} \widetilde{r}^b_{M|-}(t_b|t_s) r^s_B(t_s)
\end{aligned}
$$

$$
\begin{aligned}
f^b_{-+}(t_b,t_s) &= \pi_{BG} r^b_B(t_b)\widetilde{r}^s_{G|-}(t_s|t_b)\\
f^b_{-0}(t_b,t_s) &= \pi_{BM} r^b_B(t_b)\widetilde{r}^s_{M|-}(t_s|t_b)\\
f^b_{--}(t_b,t_s) &= \pi_{BB} r^b_B(t_b)\widetilde{r}^s_{B|-}(t_s|t_b)\\
f^s_{-+}(t_b,t_s) &= \pi_{BG} \widetilde{r}^b_{B|+}(t_b|t_s) r^s_G(t_s)\\
f^s_{-0}(t_b,t_s) &= \pi_{BM} \widetilde{r}^b_{B|0}(t_b|t_s) r^s_M(t_s)\\
f^s_{--}(t_b,t_s) &= \pi_{BB} \widetilde{r}^b_{B|-}(t_b|t_s) r^s_B(t_s)
\end{aligned}
$$

$$
\begin{aligned}
f^b_{+S}(t_b) &= \pi_{GG} r^b_G(t_b)(1-\widetilde{R}^s_{G|+}(T|t_b)) + \pi_{GM} r^b_G(t_b)(1-\widetilde{R}^s_{M|+}(T|t_b)) + \pi_{GB} r^b_G(t_b)(1-\widetilde{R}^s_{B|+}(T|t_b))\\
f^b_{0S}(t_b) &= \pi_{MG} r^b_M(t_b)(1-\widetilde{R}^s_{G|0}(T|t_b)) + \pi_{MM} r^b_M(t_b)(1-\widetilde{R}^s_{M|0}(T|t_b)) + \pi_{MB} r^b_M(t_b)(1-\widetilde{R}^s_{B|0}(T|t_b))\\
f^b_{-S}(t_b) &= \pi_{BG} r^b_B(t_b)(1-\widetilde{R}^s_{G|-}(T|t_b)) + \pi_{BM} r^b_B(t_b)(1-\widetilde{R}^s_{M|-}(T|t_b)) + \pi_{BB} r^b_B(t_b)(1-\widetilde{R}^s_{B|-}(T|t_b))\\
f^s_{S+}(t_s) &= \pi_{GG}(1-\widetilde{R}^b_{G|+}(T|t_s)) r^s_G(t_s) + \pi_{MG}(1-\widetilde{R}^b_{M|+}(T|t_s)) r^s_G(t_s) + \pi_{BG}(1-\widetilde{R}^b_{B|+}(T|t_s)) r^s_G(t_s)\\
f^s_{S0}(t_s) &= \pi_{GM}(1-\widetilde{R}^b_{G|0}(T|t_s)) r^s_M(t_s) + \pi_{MM}(1-\widetilde{R}^b_{M|0}(T|t_s)) r^s_M(t_s) + \pi_{BM}(1-\widetilde{R}^b_{B|0}(T|t_s)) r^s_M(t_s)\\
f^s_{S-}(t_s) &= \pi_{GB}(1-\widetilde{R}^b_{G|-}(T|t_s)) r^s_B(t_s) + \pi_{MB}(1-\widetilde{R}^b_{M|-}(T|t_s)) r^s_B(t_s) + \pi_{BB}(1-\widetilde{R}^b_{B|-}(T|t_s)) r^s_B(t_s)
\end{aligned}
$$

$$
\begin{aligned}
F_{SS} &= \pi_{GG}(1-R^b_G(T))(1-R^s_G(T)) + \pi_{GM}(1-R^b_G(T))(1-R^s_M(T)) + \pi_{GB}(1-R^b_G(T))(1-R^s_B(T))\\
&+ \pi_{MG}(1-R^b_M(T))(1-R^s_G(T)) + \pi_{MM}(1-R^b_M(T))(1-R^s_M(T)) + \pi_{MB}(1-R^b_M(T))(1-R^s_B(T))\\
&+ \pi_{BG}(1-R^b_B(T))(1-R^s_G(T)) + \pi_{BM}(1-R^b_B(T))(1-R^s_M(T)) + \pi_{BB}(1-R^b_B(T))(1-R^s_B(T))
\end{aligned}
$$

### Outcome Probabilities

| | |
|---|---|
| $\pi_{GG}$ | Probability of good outcome for both traders |
| $\pi_{GM}$ | Probability of good outcome for buyer, mediocre for seller |
| $\pi_{GB}$ | Probability of good outcome for buyer, bad for seller |
| $\pi_{MG}$ | Probability of mediocre outcome for buyer, good for seller |
| $\pi_{MM}$ | Probability of mediocre outcome for both traders |
| $\pi_{MB}$ | Probability of mediocre outcome for buyer, bad for seller |
| $\pi_{BG}$ | Probability of bad outcome for buyer, good for seller |
| $\pi_{BM}$ | Probability of bad outcome for buyer, mediocre for seller |
| $\pi_{BB}$ | Probability of bad outcome for both traders |

### Reporting Probabilities When Trader Rates First

| | |
|---|---|
| $\rho^b_G$ | Prob. that satisfied buyers report positive feedback if seller has not yet rated |
| $\rho^b_M$ | Prob. that mildly dissatisfied buyers report neutral feedback if seller has not yet rated |
| $\rho^b_B$ | Prob. that very dissatisfied buyers report negative feedback if seller has not yet rated |
| $\rho^s_G$ | Prob. that satisfied sellers report positive feedback if buyer has not yet rated |
| $\rho^s_M$ | Prob. that mildly dissatisfied sellers report neutral feedback if buyer has not yet rated |
| $\rho^s_B$ | Prob. that very dissatisfied sellers report negative feedback if buyer has not yet rated |

### Reporting Hazard Rate Multipliers When Trader Rates Second (Conditional on rating received by partner)

| | | | |
|---|---|---|---|
| $\alpha^b_{G\|+}$ | $\alpha^b_{G\|0}$ | $\alpha^b_{G\|-}$ | Multipliers of a satisfied buyer's hazard rate of reporting positive feedback |
| $\alpha^b_{M\|+}$ | $\alpha^b_{M\|0}$ | $\alpha^b_{M\|-}$ | Multipliers of a mildly dissatisfied buyer's hazard rate of reporting neutral feedback |
| $\alpha^b_{B\|+}$ | $\alpha^b_{B\|0}$ | $\alpha^b_{B\|-}$ | Multipliers of a very dissatisfied buyer's hazard rate of reporting negative feedback |
| $\alpha^s_{G\|+}$ | $\alpha^s_{G\|0}$ | $\alpha^s_{G\|-}$ | Multipliers of a satisfied seller's hazard rate of reporting positive feedback |
| $\alpha^s_{M\|+}$ | $\alpha^s_{M\|0}$ | $\alpha^s_{M\|-}$ | Multipliers of a mildly dissatisfied seller's hazard rate of reporting neutral feedback |
| $\alpha^s_{B\|+}$ | $\alpha^s_{B\|0}$ | $\alpha^s_{B\|-}$ | Multipliers of a very dissatisfied seller's hazard rate of reporting negative feedback |

## Appendix IV: Maple code (completes proof of Proposition 2)

```
restart;
# MAPLE code that proves the validity of the rank condition
# for any N>=2, under assumptions (A1) and (A2)
#
# Specifically, the code shows that, under (A1) and (A2)
# for any integer N the rank of the Jacobian matrix of model (1)
# is equal to the number of independent unknown variables
# Replace the following statement with the
# maximum N to test for
MAXN:=10;
# NOTATION
# unknown model variables:
# pi[i1,i2] -- probability of   latent outcome i1,i2
# rho[j,i]  -- probability that buyer  reports j when observing i
# tau[j,i]  -- probability that seller reports j when observing i
# manifest probabilities:
# F[j1,j2]  -- probability of observing feedback pattern j1,j2
# don't clutter output with evaluation details
# remove this if more detail is desired
printlevel:=-5;
# load linear algebra package
with(linalg):
# prove rank condition separately for each N
for N from 2 to MAXN do
M:=N;
```

```
# initialize all variables
# eqset, varsets will eventually contain the sets of
# independent equations and variables respectively
eqset :={};
varset:={};
# unevaluate arrays from possible prior iterations
for i from 1 to N+1 do
for j from 1 to N+1 do
  rho[j,i]:= evaln(rho[j,i]);
  tau[j,i]:= evaln(tau[j,i]);
  if i<N+1 and j<N+1 then pi[i,j]:=evaln(pi[i,j]); end if;
end do;
end do;
# garbage collect
gc();
# we assume truthful reporting;
# thus, set all "untruthful" reporting probabilities to zero
# ("untruthful" reporting means observe i but report j<>i)
for i from 1 to N do
for j from 1 to N do
  if i<>j then rho[j,i]:=0; tau[j,i]:=0; end if;
end do;
end do;
# set prob. of remaining silent = 1 - sum of prob. of reporting anything else
for i1 from 1 to N do
  rho[M+1,i1]:=1;
  tau[M+1,i1]:=1;
```

```
  for j1 from 1 to M do
    rho[M+1,i1]:=rho[M+1,i1]-rho[j1,i1];
    tau[M+1,i1]:=tau[M+1,i1]-tau[j1,i1];
  end do;
end do;
# set probability of outcome #N = 1 - sum of prob. of outcomes 1..N-1
pi[N,N]:=1;
for i1 from 1 to N do
for i2 from 1 to N do
  if i1+i2<2*N then pi[N,N]:=pi[N,N]-pi[i1,i2]; end if;
end do;
end do;
# now generate the model's equations
# store equations in eqset
# store variables in varset
for j1 from 1 to M+1 do
for j2 from 1 to M+1 do
  if j1+j2<2*M+2 then
    eq[j1,j2]:=F[j1,j2];
    for i1 from 1 to N do
    for i2 from 1 to N do
         eq[j1,j2]:=eq[j1,j2]-pi[i1,i2]*rho[j1,i1]*tau[j2,i2];
         if i1+i2<2*N then varset := varset union {pi[i1,i2]}; end if;
         if rho[j1,i1]<>0 and j1 < M+1 then varset:= varset union {rho[j1,i1]}; end if;
         if tau[j2,i2]<>0 and j2 < M+1 then varset:= varset union {tau[j2,i2]}; end if;
      end do;
    end do;
```

```
    eqset := eqset union {eq[j1,j2]};

 end if;

end do;

end do;

# convert sets to lists (vectors) so that Jacobian can be calculated

equl:=convert(eqset,list);

varl:=convert(varset,list);

# calculate the rank of the jacobian of the vector of equations

# with respect to the unknown variables

jac1:=jacobian(equl,varl);

rank1:=rank(jac1);

printf("# of possible outcomes N = %2d ::

        Jacobian rank %3d equals number of unknown variables %3d\n",

        N, rank1, nops(varl));

end do;
```