Last Minute Feedback*

Tobias J. Klein
Christian Lambertz
Giancarlo Spagnolo
Konrad O. Stahl†

first version: July 9, 2005
this version: January 2, 2007

Abstract

Feedback mechanisms in electronic markets allow trading partners to rate each other after a transaction and to observe past ratings given by other market participants. Such mechanisms are considered crucial for the success of internet trading platforms, where alternative governance methods are often missing or too expensive. However, the concern has been raised that on eBay positive feedbacks may be given in order to encourage reciprocation, and negative feedbacks might not be given because of the fear of retaliative negative feedback. If quantitatively important, these practices substantially reduce the informational content of reputation indicators. In this paper, we investigate the micro structure of the eBay feedback mechanism in light of an analysis of the sequencing and timing of feedbacks. Since the transactions *per se* are unobservable, we derive bounds on probabilities of buyer and seller satisfaction by exploiting alternative sets of identifying assumptions. We argue that substantive bias is potentially accumulated in the reputation measures and conclude with easy-to-implement suggestions geared at improving the design of the feedback mechanism in terms of the informational content of its records.

JEL Classification: D44, L15, L86.
Keywords: eBay, reputation mechanism, strategic feedback behavior, informational content, reciprocity, fear of retaliation.

*This paper has been entirely rewritten in the second half of 2006. We are grateful to Florian Hauber and Johannes Koenen for excellent research assistance. We truly appreciate having had the opportunity to extensively work on this paper during the 2005 C.E.P.R. European Summer Symposium in Economic Theory in Gerzensee, Switzerland. We thank Heski Bar-Isaac, Florian Müller, Axel Ockenfels, Klaus Schmidt, Henry Schneider, Mike Ward, and two anonymous referees for helpful comments and audiences at the 2006 SFB Conference in Caputh, at the 2006 International IO Conference in Boston, at the 2006 Meeting of the European Economic Association in Vienna, at the 2006 European Association for Research in Industrial Economic (EARIE) conference in Amsterdam, at the 2006 ZEW Internet Policy Workshop in Mannheim and in seminars in Berlin, Mannheim, Munich and Stockholm for lively discussions.

†T. Klein (klein@econ.uni-mannheim.de) and C. Lambertz (lambertz@econ.uni-mannheim.de): University of Mannheim; K. Stahl (kos@econ.uni-mannheim.de): University of Mannheim, C.E.P.R., *CES*ifo, and ZEW; G. Spagnolo (giancaspagnolo@yahoo.com): Consip, Stockholm School of Economics, and C.E.P.R. *Address:* University of Mannheim, Department of Economics, 68131 Mannheim, Germany.

1 Introduction

Reputation is a fundamental concept in economics, and a main governance mechanism for most real world market transactions, active whether or not other governance methods—like explicit contracts—are present. In the theoretical literature on reputation, it is typically assumed that either the reputation bearer's behavior is publicly observed by potential future partners, or that privately observed behavior is truthfully communicated from one to another market participant (e.g. through "word of mouth communication").[1] However, empirical work on how reputation is built and transferred in the market is missing, as word of mouth communication is rather hard to register and analyze. Electronic markets, with their publicly observable feedback records, offer economists the important chance to empirically analyze the way reputation forms and acts. In this paper we take this chance and analyze the micro structure of the reputation mechanism used on eBay, the largest market ever in terms of number of participants. In this and other electronic trading platforms, the reputation of an agent consists of the ratings of her trading partners in past transactions. The feedback mechanism is claimed to be crucial for its success as well the success of other trading platforms, as trading on these platforms goes along with a particularly wide room for opportunistic behavior on both sides of the market. Anonymity and distance allow sellers to cheat on the quality of the good. Likewise, buyers can be dishonest concerning their payment behavior.[2] As long as feedbacks truthfully reflect the trading partners' experiences in the transactions with the agent in question, the accumulated feedbacks indeed provide valuable information for potential trading partners. Moreover, the feedback mechanism thereby disciplines the aforementioned forms of opportunism with the threat that, if that agent misbehaves today, she will get a bad feedback and will therefore be bypassed by other traders in the future. Public statements by eBay emphasize this (desired) effect of the feedback mechanism.[3]

The central question addressed in this paper is whether the feedback mechanism implemented by eBay provides incentives to report experiences truthfully. Rather than focusing on the *effects* of reputation, for example on prices or the probability of selling, we focus on the *design* of the reputation mechanism and its impact on the *reputation building* process.[4]

[1] See, e.g., Livingston and Evans (2004) and Cabral and Hortaçsu (2006) for eBay.

[2] According to the Internet Crime Complaint Center (IC3) 2005 Internet Fraud Crime Report "internet auction fraud was by far the most reported offense, comprising 62.7% of [97,076] referred complaints." See http://www.ic3.gov/media/annualreports.aspx (September 2006). Likewise, the FTC reports that "internet auction fraud is on the rise, with an increasing number of consumers complaining about sellers who deliver their advertised goods late or not at all, or deliver something far less valuable than promised." See the FTC's "Top Ten Dot Cons" on http://www.ftc.gov/bcp/conline/edcams/dotcon/auction.htm (February 2006).

[3] eBay states that the feedback "comments and ratings are valuable indicators of your reputation as a buyer or seller on eBay," see http://pages.ebay.com/help/feedback/questions/feedback.html (February 2006). Moreover, in the founder's letter posted on February 26, 1996, Pierre Omidyar claims that "some people are dishonest. Or deceptive...But here, those people can't hide. We'll drive them away." See http://pages.ebay.com/services/forum/feedback-foundersnote.html (February 2006).

[4] See Dellarocas (2005) for a useful survey of recent research on reputation mechanisms. The effects of seller reputation on prices and the probability of selling the object are usually found to be negligible or positive. See, for example, Melnik and Alm (2002), Bajari and Hortaçsu (2003), Cabral and Hortaçsu (2006), Livingston and Evans (2004), Lucking-Reiley, Bryan, Prasad, and Reeves (2005), Houser and Wooders (forthcoming). See also Bajari and Hortaçsu (2004) as well as Resnick, Zeckhauser, Swanson, and Lockwood

On eBay, both the seller and the buyer of an object are allowed to rate each other after a transaction. It has been argued that leaving a positive feedback might be driven by expectations on *feedback reciprocity* which induces the trading partner to post a positive mark in return.[5] Conversely, agents dissatisfied with their trading partners might refrain from leaving negative feedbacks at all, as they anticipate revenge. Consequently, if agents who give the first feedback expect the opponent to reciprocate positive, or retaliate negative feedbacks, then, relative to truthful reports, one would expect negative first feedbacks to be rare and positive ones to be common, a pattern that is usually found.[6] If an eBay user is not fully aware of this fact, she might then overestimate the informational content of feedback records.[7]

Whereas eBay does not document the trading process that follows an auction on its platform, the traders' reactions to their trading partners' behavior in form of feedbacks are observable. From these, we cannot directly infer the joint distribution of buyer and seller satisfaction in order to evaluate their truthfulness, but must resort to theory in order to interpret the evidence. Naturally, that must be based on the mechanism currently implemented by eBay. Key elements of this are the following:

1. Feedbacks are immediately observable.
2. Feedbacks, once given, cannot be modified.
3. Feedbacks can be withdrawn upon the mutual consent of both trading partners.
4. The feedback period is of stochastic length and trading partners can always react to a feedback.

In principle, each of these elements can generate opportunistic behavior. 1. implies that a user risks a retaliative negative once he leaves a negative feedback. A consequence of 2. is that the first feedback giver is locked in and is thereby put at the mercy of his trading partner. If she behaves opportunistically towards establishing a good feedback record, a first feedback giver will therefore not give a negative feedback even if it is justified. Instead, she might wait for her partner to give the first feedback. If both decide to do so, no feedback is left at all. 3. implies

(2004) for an overview.

[5]This tendency to reciprocate may be due to behavioral components in agents' decision making processes, similar to the ones found by Fehr and Schmidt (1999); e.g. the attempt to build up a reputation as a "reciprocator" or "impersonator" in order to discourage future negative feedbacks and encourage positive ones—"the high courtesy equilibrium" of Resnick and Zeckhauser (2002)—, or the combination of both motives. Dellarocas and Wood (2006) relate the motivation for leaving positive feedback to the user's expectation of reciprocal behavior from their trading partners.

[6]See, for example, Resnick, Kuwabara, Zeckhauser, and Friedman (2000), Resnick and Zeckhauser (2002), Bajari and Hortaçsu (2004), Cabral and Hortaçsu (2006), and Chwelos and Dhar (2005). If quantitatively important, both forms of strategic behavior reduce the informational content of the reputation index and in particular bias it towards positive evaluations of the transaction This point is also made by a very recent study of Dellarocas and Wood (2006) estimating that only about 86 per cent of eBay users are actually happy with the underlying transaction.

[7]Jin and Kato (2002) find in a field experiment that "at least some buyers" overestimate the informational content of feedback score and "drastically underestimate the risk of trading online." Likewise, Resnick, Zeckhauser, Swanson, and Lockwood (2004) question whether price premia, which they find, reflect a reputation equilibrium, and should in fact not be observed in the data.

that the second feedback giver has an incentive to use a negative feedback as a strategic device in order to induce her trading partner to agree to a mutual withdrawal of negative feedbacks. She might consider doing so even if she is in fact happy with the behavior of her trading partner in the underlying transaction.

In the empirical analysis, we describe the relationship between the timing of feedbacks, their type (positive, neutral, or negative), the rating sequence (who rates first), and trading partners' experience on eBay. We document evidence suggesting that rating behavior on eBay is in fact led by strategic considerations and derive bounds on buyer and seller satisfaction, stressing what can be learned from observed feedbacks under varying identifying assumptions, and what cannot be learned without making stronger assumptions.

In our discussion of the institutional details we discuss the fourth point mentioned above, namely the "ending rule" of the period in which trading partners can post their feedback ratings, in more detail. In particular, it is important whether this period is of fixed or stochastic length. Sincerely dissatisfied participants may be (rightly) concerned that if they post a justified negative feedback the trading partner could retaliate with a non-justified negative. An analysis of the institutional details reveals that, literally, eBay "only guarantees" that a feedback can be left within 90 days after the auction ends. Whereas this does not exclude the possibility that feedback can be left thereafter, a possible perception of this somewhat opaque statement by parts of the user community is that the end of the 90th day constitutes a final, deterministic deadline. If this was the case negative feedback could be left in the "last minute" to minimize the risk of retaliation. This could induce truly dissatisfied users to strategically leave negative feedbacks more often: if the negative is left in the last minute, then, the trading partner will not have the time to retaliate. We will refer to those feedbacks as "last minute feedbacks" and denote them by the acronym LMF.

Such a strategic situation resembles closely that leading to "last minute bidding" in English auctions with fixed ending time (Roth and Ockenfels, 2002). In both cases a "last minute action" is exploited in order to prevent the opponents' reaction. However, if one were to consider mechanisms without fixed ending times, agents in an auction would still prefer placing a bid to abstaining.[8] On the contrary, giving a negative feedback becomes less attractive because of the fear of retaliation. Therefore, from a welfare point of view, the presence of a last minute is desirable in the context of feedbacks, whereas in the context of bids, it is not necessarily so.[9]

In light of these considerations, it is surprising that the end of the feedback period is in fact stochastic on eBay. Even more surprisingly, by investigating the institutional details we found that once a first feedback is left, the trading partner always has at least an additional 90 days to reply with a feedback! That is, the real ending rule of the feedback period on eBay is similar to the one of Amazon for auctions.[10] With its similar structure, the eBay feedback mechanism may thus discourage truthful negative ratings by the potential first feedback giver, by giving the opponent enough time to retaliate. This opens the questions why eBay chose such a peculiar feedback mechanism, and why eBay is so opaque when describing its functioning.

We conclude suggesting some simple ways to discourage opportunistic first feedback giv-

[8]For example, Amazon type auctions end only when no more bids are placed.

[9]This is the case if truthful reports are considered welfare optimal.

[10]There, a "last minute" bid prolongs the auction period automatically. See http://www.amazon.com/gp/help/customer/display.html?ie=UTF8&nodeId=1161360 (September 2006).

ing and to improve on agents' incentives to express their possible dissatisfaction. An obvious improvement would be to actually introduce the absent deterministic last minute wrongly perceived by some eBay users. It would be even better to combine this with feedback revelation to the trading partner and the community only after no more feedbacks can be left. Finally, mutual feedback withdrawal should not be allowed for.

In Section 2, we describe the feedback mechanism in detail. Section 3 contains an analysis of feedback behavior with a focus on the timing of ratings. We derive simple bounds for buyer and seller satisfaction in Section 4. Section 5 discusses the costs and benefits of possible improvements of such bilateral feedback mechanisms.

2 The eBay Feedback Mechanism

eBay is by far the biggest internet trading platform. It brings together both private and professional buyers and sellers. In 2005, the number of listings exceeded 1.9 billion and eBay's gross merchandise volume accounted to more than 44 billion U.S. dollars.[11] Amongst other services, eBay provides a second price auction mechanism in which the seller describes the object and specifies a reservation price as well as the length of the auction period. Then, potential buyers can enter their bids.

When an item is sold, eBay sends a notification to the seller and the buyer who should then contact each other and arrange payment and shipping. As an intermediary, eBay assumes no responsibility in the process after the auction has ended and is only involved if a dispute arises. After the auction, both the seller and the buyer can leave feedback on their trading partner. Neither party is required to leave feedback, but is actively encouraged by eBay to do so. Resnick and Zeckhauser (2002) who have access to auction and feedback data report that buyers commented on sellers 52.1% of the time and sellers on buyers 60.6% of the time. A feedback is a "positive," "neutral," or "negative" rating accompanied by a textual comment. For ease of the exposition, we follow the literature and occasionally group neutral and negative marks together, see e.g. Resnick and Zeckhauser (2002) and Cabral and Hortaçsu (2006). We will refer to them as negatives.[12]

As already mentioned, on eBay, feedbacks are immediately observable to the counterpart. For every user, eBay keeps a feedback record which contains all feedbacks received and left from transactions in which she was involved.[13] A recorded feedback cannot be unilaterally removed. However, as at least some observers appear to have noticed feedbacks *can* be withdrawn if both parties agree to.[14] All marks a member has ever received are summarized in a feedback

[11] See http://investor.ebay.com/news/Q405/EBAY0118-123321.pdf (February 2006).

[12] However, we should emphasize that separating neutral from negative marks would not qualitatively change our results.

[13] eBay also offers internet shop services. Thus, feedbacks may also be based on experiences in trading via this channel, rather than auction trading.

[14] eBay states that "[a]fter both parties have agreed to withdraw the feedback, both parties will have their feedback scores adjusted at the same time... eBay will add a note to the feedback comment, saying that the feedback was mutually withdrawn... If you haven't left feedback for your trading partner and you go through the Mutual Feedback Withdrawal process, you will no longer be able to leave feedback for that transaction... You may only request Mutual Feedback Withdrawal once for every feedback left... Members may initiate a request to mutually

score and several summary statistics including the percentage of positive feedbacks.[15]

Given this mechanism, actual feedback behavior can be influenced by strategic considerations, arising whenever agents anticipate the opponents' reactions when giving feedback. As opposed to truthful reports this yields potentially biased reports, as they are not only influenced by the transaction outcome. Strategic considerations are important if a user is planning to interact with other users in the future, and therefore attaches positive value to her own reputation. That is, she derives positive utility (expected payoff) from a positive and negative utility from a negative feedback received. This will be the case as long as there is some potential future trading partner believing that the feedback score is informative about the likely behavior of its holder.[16]

Feedbacks can also be strategically given in order to influence the outcome of the current transaction. However, as we cannot observe this, we abstain from this consideration. Indeed, this strategic influence may yield to both, efficiency improving as well as efficiency decreasing outcomes of the transaction. [17]

At any rate, on eBay, the overwhelming share of all feedback is positive. Resnick and Zeckhauser (2002) report that, in their data, ratings left by sellers and buyers were positive in 99.1 and 98.1 per cent of the cases, respectively. From another point of view, the average seller's share of negative ratings in the study by Cabral and Hortaçsu (2006) is 0.9 per cent of all comments. This study confirms these scales as in our data about 99 per cent of the observed feedbacks are positive.

If reporting was always truthful and missing feedback would be missing at random these numbers would suggest that the overwhelming share of transactions on eBay simply results in contented trading partners.[18] On the other hand, there are indications that these numbers are biased; more precisely, that users are reluctant to leave negative feedback because of fear of retaliation by the other party. Such fear of retaliatory negative feedback is often expressed in discussion boards and newsgroups on eBay. Statements similar to the following from a buyer can be found on many forums[19]:

> Just last week, I had my first unpleasant experience in five years of eBay'ing. I received an item from a seller who had not left feedback for me (I mailed my money

withdraw feedback within 30 days of either person leaving feedback or within 90 days of the transaction end date, whichever is later." Taken from http://pages.ebay.com/help/feedback/questions/mutual-withdrawal.html (September 2006).

[15]The feedback score is calculated as the number of users who left positive feedback minus the number of users who left negative feedback.

[16]As was already pointed out in footnote 4 price effects of reputation are usually found to be nonnegative. Therefore, a "good" reputation on eBay is currently valuable to sellers. In principle, potential buyers in an auction could distinguish feedbacks the seller has received as a seller from feedbacks she has received as a buyer. However, it is a complex task to infer separate summary statistics from the records. See also Cabral and Hortaçsu (2006) who find that at least some sellers were able to build up their reputation as buyers. Even pure buyers can benefit from a "good" reputation record since sellers are allowed to exclude buyers from their auctions. This is possible on the basis of their subjective judgement of a bidder's reputation record.

[17]For example, a negative feedback giver may, by offering to withdraw that feedback, modify the outcome of the transaction by thus generating a hold up, or by removing a hold up brought about by his trading partner.

[18]Dellarocas and Wood (2006) investigate the consequences of non-random missing feedbacks on feedback scores. They find that dissatisfied traders are more likely not to give a feedback.

[19]Quotes are taken from http://ideas.4brad.com/archives/000018.html (February 2006).

> order the day after the auction ended). I was not happy with the item - flaws were not disclosed in the listing - and I notified the seller. After three e-mails and three phone calls went unanswered, I left negative feedback for her. She turned around and posted retaliatory negative feedback for me ruining my 100% rating. Indeed, the system needs to be improved.

Other users do not expose themselves to this risk

> In the past I've not left any neg[ative] feedback as I'm afraid of revenge feedback that'll paint me as a bad trading partner... the dodgy seller ends up with getting away with it just to rip someone else off.

Another user notes

> As a buyer I have had problems with false item descriptions, even if you get a refund ... you end up paying postage for the item to you and back. Up till now I have not left any feedback for these bastards because of revenge.

and one concludes that

> I have been basing my purchase decisions [on eBay] on sellers' feedback scores. I had no idea these scores are so unreliable ... They are holding this feedback system out as the reason we should trust sellers, but the system has little to no basis in truth ... I suspect there are many, many people out there who have had actual monetary losses from this behavior.

Another user rounds up the dilemma[20]:

> Sooner or later we all face this dilemma on e-Bay. Do we slag an obvious jerk with a negative feedback, only to get a retaliatory negative feedback from him. You have to decide if it's worth it. Always check out his feedback first. See if she posts retaliatory feedbacks. Avoid him like the plague if she does. In your case, seeing as how you aren't out any cash, I would just let this one slide. Let this moron fester in his own little crooked world. There are a lot of goofs out there in e-Bayland, just steer clear of them if possible. IMHO [in my humble opinion], save your negative feedbacks for the really bad experiences that cost you serious money. Cheers!

Such debates show that users are aware of the risk of negative feedback retaliation, or revenge. Retaliating against deserved negative feedbacks (and reciprocating positives) may be useful, for example, to build a reputation of being an imitator, who always replies strategically to a positive feedback with a positive one, and to a negative feedback with a negative one. Such a reputation may be valuable because it encourages future partners to give positive feedbacks and discourages them from giving negative ones. eBay even sells a service to sellers allowing them to automatically reciprocate positive feedbacks.[21] Such behavior is in principle observable to other users on eBay.[22] The expectation that a deserved negative feedback may induce a

[20] Taken from http://antiqueradios.com/forums/Forum14/HTML/000994.html (February 2006).

[21] The price for an online seller tool which includes this service is currently $15.99 a month, see http://pages.ebay.com/sell/automation.html (February 2006) for a description.

[22] In particular, the feedbacks a user gets and the replies she leaves can be inferred from her feedback record.

non-deserved retaliation is further justified by the possibility to withdraw feedbacks by mutual consent. If a party receives a negative, it can retaliate with another negative to have "something to trade" to persuade its partner to withdraw.[23]

To alleviate concerns regarding retaliatory negative feedback, it is an often-repeated suggestion to "snipe" negative feedback, i.e. leaving it so late that the other party cannot react. The idea is explained by this discussant[24]:

> The secret... is to wait until the 90 day feedback period is nearly up and then zap em w[ith a] negative feedback when they only have a few hours remaining to respond... That way they can't retaliate... This only wor[ks] if you are able to hold a grudge for 90 days...

Moreover, it has been suggested in various newsgroups to set up a service that automates strategic feedback timing. In a typical conversation, a user suggests[25]

> will someone out there please invent FEEDBACK SNIPER SOFTWARE that allows one to leave feedback (good or bad) at the last second? that way, you can leave legit[imate] bad feedback w[ith] no fear of retaliatory bad feedback left for you- thus purifying the ebay world, making ebay stock go up, and just making ebay a better community as a whole. i do not leave deserved bad feedback for fear of retaliatory bad feedback left on me!!!

And indeed, Auctionhawk, a company specialized on offering services around eBay, developed and advertised a service, for payment, to give feedback in the last minute.[26]

However, "feedback sniping" is not a safe strategy. While eBay *guarantees* that feedback comments are recorded if left within 90 days after the end of the auction, this does not exclude the possibility of leaving feedback after this 90 day period. In eBay's own, opaque words: "eBay only commits to items being available for 90 days, so if it is greater than 90 days you may not be able to leave feedback."[27]

After 90 days, eBay removes the link on a member's personal "My eBay" page that encourages one to leave feedback. However, since the item number identifying a particular transaction is known in principle, one might still be able to leave feedback for a transaction by doing so manually. At some stage, around the 120th after the end of the auction, it becomes impossible

[23] In game-theoretic jargon, the possibility to withdraw feedbacks by mutual consent makes retaliating a dominant strategy in the subgame starting after a first negative feedback is posted, see Figure 1 and the accompanying explanation.

[24] See http://www.the-gas-station.com/messages.cfm?type=normal&thread_id=49933&lastdays=2000& (February 2006).

[25] See, e.g., http://community.auctionsniper.com/groupee/forums/a/tpc/f/785608021/m/308108399/r/3721016131 (February 2006). The quotes that follow are taken from this site.

[26] See http://auctionbytes.com/cab/abn/y04/m08/i10/s01 (February 2006). A free reminder service for "last minute feedback" is offered by U.K. Auction Watch at http://www.ukauctionhelp.co.uk/remindme.php (February 2006).

[27] See http://pages.ebay.com/help/feedback/questions/leaving-feedback.html (September 2006).

to leave a (first) feedback.[28] Once a rating is left, though, it is recorded in the partners' profiles as "feedback received" or "feedback left," respectively. In addition, the corresponding auction number is reported and linked to the auction details for at least another 90 days during which a second feedback can be submitted. In consequence, a trading partner (who has not yet submitted a rating himself) will always have at least 90 days to reciprocate or retaliate, independently of how early or late after the end of an auction feedback was left for him. This makes successful "sniping" of negative feedback technically impossible: There is no way to leave a feedback without giving the opponent the possibility to react.

Nevertheless, the newsgroup discussions show that at least some eBay users are not aware of this, and perceive the 90 day deadline as a fixed, deterministic end after which no more (e.g. retaliatory) feedback can be left. Accordingly, McDonald and Slawson (2002) note that "some users attempt to avoid retaliatory negative feedback by 'sniping a negative' ".

For many eBay users, therefore, the perceived structure of the feedback period is the one depicted in the state chart in Figure 1. Each circle in the graph represents a state and each arrow a transition from a state into another state or into itself. Such a transition happens in every instant of time. For example, we enter the feedback game from the left. Then, we are in the state in which nobody has left feedback so that it can still be left by both. In the next instant of time, either no one rates and the feedback period is not over, or only one of the two parties rates, or both rate each other simultaneously, or the feedback period is over. Depending on the actions of the players we transit into another state.

The last (grey shaded) state is always the payoff state. Payoffs are to be understood as expected payoffs from future transactions via eBay, conditional upon the feedback outcome. Depending on the history, either no feedbacks have been left, or feedbacks have been left without being withdrawn—the usual case—, or feedbacks have been left and withdrawn thereafter. Note that here, for ease of the exposition, we make the simplifying assumption that we can always enter the feedback withdrawal process after at least one feedback has been left. In reality, every player may initiate this only once, see footnote 14 for details. Note that only a subset of the users on eBay is likely to be aware of the possibility of withdrawal.

The dashed part of the graph represent the misperception of the existence of a last minute. To be more specific, the misperception is that after 90 days there is a transition into a "last minute" state in which the trading partner cannot react to a feedback. Most importantly, the chart (without the dashed part) indicates that once a first feedback is left, in fact, the trading partner has *always* the opportunity to react with a second feedback.

In the next section we first document feedback behavior on eBay, focussing on its timing, and investigate whether the fear of retaliation and the misperception of the feedback game as described above indeed lead to "last minute" feedback, and if so, amongst which players.

3 Descriptive Evidence

The data for the empirical analysis were collected in the second quarter of 2005 from the eBay platform. Starting from randomly drawn users we created a data set consisting of $2,470,522$

[28] We learned this from our data.

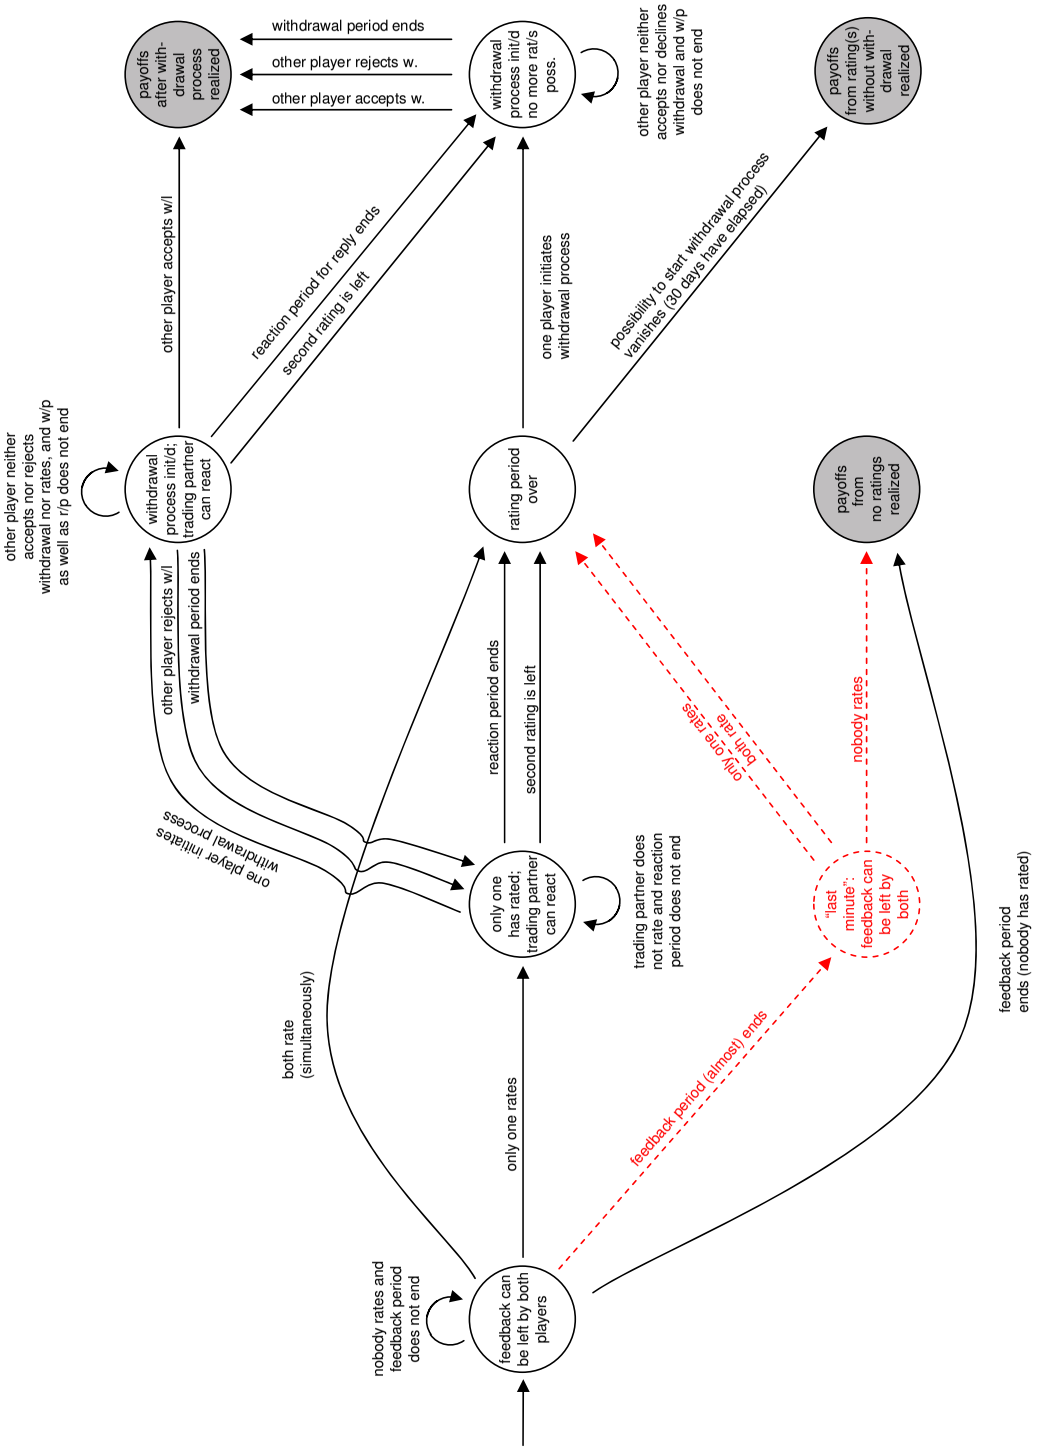

Figure 1: State chart.

auction records including respective feedbacks and their timing. By construction, the data include auctions for which at least one feedback was left. It is a random sample with respect to the category of the auctioned good which we think is appropriate for the purpose of this empirical analysis since we want to study feedback behavior *in general*. The data collection procedure is described in more detail in Appendix A. Additional Tables and Figures are contained in Appendix B. In the sequel, we occasionally refer to the first and second feedback as *feedback* and *reply*, respectively.

Table 1 contains summary statistics. It shows that some first feedbacks are given immediately after the end of the auction and others are given as late as on the 146th day. Seller feedback scores are substantially larger than buyer scores, indicating that sellers are active much more frequently in eBay, and thus more experienced than buyers. There is a considerable number of big players in our data set: one professional seller has a feedback scores which is as high as $236,400$. Finally, observe that the first feedback is given more often by the less experienced player and more often by buyers. This is the first aspect we analyze in more detail.

3.1 Rating Sequence

As we have already pointed out, a user who gives feedback first is locked in and put at the mercy of her trading partner. At the same time, going first with a positive mark could induce a reciprocative positive feedback. The mechanism implemented by eBay thereby gives rise to strong incentives for users to rate first if they believe that their trading partner is satisfied with their performance in the transaction. On eBay, typically, payment precedes shipment. Thus, an opportunistic buyer who has paid correctly and observes that his payment has arrived at the seller can remove a potential hold up against the seller by a first positive feedback in the expectation to receive a positive reply. This may induce bias in the reputation record in the case the seller did not perform well.

In Table 2, we report on the rating sequence, and in particular on the impact of experience. We define a user to be experienced if his feedback score is at least of size 20.[29] Increased experience clearly results in improved understanding of the mechanism prescribed by eBay, and of the strategic situation the trader finds heself in. Therefore, we expect strategic behavior to be more pronounced for experienced players. At the same time, experience might go along with smoother transactions on eBay as it allows users to learn from past transactions and improve on their performance, e.g. when writing the auction listing, paying, or sending the good.

The table shows that if a seller is more experienced, he will be less likely to rate first, whether or not the buyer is experienced. Conversely, as buyer experience increases, *ceteris paribus*, the likelihood increases that the buyer rates first. Hence while sellers seem to wait with their rating until the buyer has posted her rating in order to not expose themselves to hold up, experienced buyers tend to actively remove that hold up by rating first. The pure sequence of rating may reflect the sequence in which the moves are conducted during the transaction. Also it is more difficult for sellers to form beliefs on how the buyer evaluates their performance, as they might feel that the good is not as described, e.g., so that they choose to wait for a positive buyer rating

[29] A cutoff value of 20 for the feedback score means that the user has participated in about 40 transactions if she receives feedback in 50 per cent of the cases. By then, we believe, she should be familiar with the rules on eBay. We tried other cutoff values which yielded similar results.

	observations	mean	std.	min.	max.
timing of first feedback in days after end of auction	2,470,522	15.93	16.10	0.00	145.29
timing of second feedback in days after end of auction	1,741,571	20.05	16.77	0.00	139.90
time difference between feedback and reply in days	1,741,571	6.16	11.62	0.00	120.23
feedback score of first feedback giver	2,458,429	11,392.68	32,243.50	-7.00	236,400.00
feedback score of second feedback giver	2,461,829	16,923.78	32,345.59	-8.00	233,143.00
seller feedback score	2,470,217	27,924.11	41,367.75	-3.00	236,400.00
buyer feedback score	2,450,041	282.82	928.63	-8.00	91,676.00
indicator for first feedback giver having higher feedback score	2,458,429	0.37	0.48	0.00	1.00
indicator for first feedback being left by seller	2,470,522	0.37	0.48	0.00	1.00

Table 1: Summary statistics.

buyer experienced	seller experienced	first feedback by seller	observations
		0.372	2,470,522
no	no	0.681	94
yes	no	0.529	899
no	yes	0.406	507,425
yes	yes	0.363	1,962,104

Experience is defined by a feedback score of at least 20.

Table 2: The effect of experience on the probability that the first feedback is given by the seller.

	(1)	(2)	(3)	(4)
seller experienced	-0.177**		-0.165**	-0.167**
	(0.016)		(0.016)	(0.016)
buyer experienced	-0.043**		-0.042**	-0.043**
	(0.001)		(0.001)	(0.001)
time of first feedback		-0.010**	-0.010**	-0.010**
		(0.000)	(0.000)	(0.000)
first feedback negative				-0.036**
				(0.004)
observations	2,470,522	2,470,522	2,470,522	2,470,522

Marginal Effects. Standard errors in parentheses. * significant at 5%; ** significant at 1%. The dependent variable is an indicator which takes on the value 1 if a first feedback is left by the seller. Experience is defined by a feedback score of at least 20.

Table 3: Probit model for first feedback given by seller.

until they rate themselves. Note, however, that this does not explain that the probability of leaving first feedback increases in their experience. Indeed this suggests opportunistic behavior on their side. Keep in mind that experienced buyers are agents who are more frequently active in the market and thus more highly value positive feedbacks.

These findings are confirmed by the probit estimates reported in Table 3. They show that seller experience has a quantitatively higher impact on the probability of the seller rating first, decreasing it by about 17 percentage points as compared to buyer experience which decreases the probability by about 4 percentage points. Additionally, these estimates show that the later a first feedback is given the more likely it is that it is given by the buyer. This seems natural as typically, on eBay, the seller only dispatches the good once she has received payment. Therefore, out of the transaction sequence, the buyer knows later on average whether the seller behaved well. Finally, our estimates indicate that if a first feedback is negative, it is less likely to be given by the seller. This could reflect the fact that buyers have more to complain about on average as, arguably, there is more room for opportunistic behavior on the side of the sellers.

In summary, the evidence on the rating sequence coupled with agents' experience shows that buyers tend to rate first the more frequently the more experienced they are. This can be interpreted as evidence for opportunistic behavior on the part of the experienced buyers.

3.2 Feedback Patterns

Table 9 in Appendix B shows feedback patterns differentiated by type of feedback provided and role of the first feedback giver in the transaction. It again demonstrates very clearly that the buyer tends to rate first and positively. Typically, the seller reacts with a positive or no mark. This table will be used later on to derive simple bounds of buyer and seller satisfaction. We will then exploit the fact that some ratings, e.g. second negatives that follow a first positive, can safely be assumed to be truthful, while other ratings such as first positives might be given out of opportunistic motives.[30]

Table 4 contains sample probabilities for the first feedback being positive, neutral, or negative. It also shows the dependence of the likelihood to observe a particular type of first feedback on experience and the role of the trading partner giving the first feedback. Column (1) shows that overall, first ratings are more likely to be negative if they are given by buyers. This corresponds well to the findings in column (4) of Table 3 and the accompanying argument.

The table also shows that the percentage of observed positive feedbacks increases in the experience of the trading partners. For buyers, compare the respective entries in columns (2) and (3), and columns (4) and (5), respectively. For sellers, compare columns (2) and (4), and (3) and (5), respectively.

This pattern could have at least two causes. First, more experience on either side could simply result in a smoother transaction and hence leave both trading partners happier on average. Second, experience could promote strategic behavior so that we observe fewer truthful negative feedbacks by experienced users. In favor of the latter we observe that the percentage of positive feedbacks that a first rating buyer leaves when evaluating the performance of the experienced seller is 0.946 if she is not experienced and 0.983 if she is. Similarly, an inexperienced seller will leave fewer positive feedbacks for an experienced buyer, given that he rates first, as compared to an experienced seller. Assuming that the truthful evaluation of a trading partner does not depend on own experience, this suggests that experience promotes strategic behavior as negative first feedbacks are observed more often by inexperienced players.[31]

Table 5 contains unconditional and conditional sample probabilities for the reply being positive, neutral, negative, or missing, with the rows indicating the conditions. In 70.5 per cent of the cases a reply is left if a feedback is left. In 70.9 per cent of the cases we observe that a positive feedback is reciprocated whereas only in 32.6 per cent of the cases a negative feedback is retaliated. This table could reflect that a first negative is only given if it is relatively unlikely that it is retaliated, whereas a positive feedback is often associated with a (desired) positive reply. This suggests that feedback scores are biased as the missing feedback (transactions in which no feedback is left) are not missing at random since relatively more negative feedbacks are missing.

[30] Observe that *a priori* second negatives are not necessarily truthful, as the feedback giver could generate a hold up by this mark. This is likely to having been the case if feedbacks were withdrawn thereafter. As we exclude mutually withdrawn feedbacks from our analysis, we therefore make this assumption only on second negatives that have not been withdrawn.

[31] Obviously, the decision of who rates first is endogenous. Hence this observation, while suggestive, does not constitute a proof as it is conditional on the rating sequence.

		(1)	(2)	(3)	(4)	(5)
	buyer experienced		no	yes	no	yes
	seller experienced		no	no	yes	yes
feedback is	positive	0.983	0.819	0.959	0.961	0.989
	neutral	0.008	0.032	0.008	0.017	0.006
	negative	0.009	0.149	0.033	0.022	0.005
	observations	2,470,522	94	899	507,425	1,962,104
buyer rates first and feedback is	positive	0.976	0.833	0.934	0.946	0.983
	neutral	0.013	0.100	0.012	0.029	0.009
	negative	0.011	0.067	0.054	0.025	0.008
	observations	1,551,582	30	423	301,385	1,249,744
seller rates first and feedback is	positive	0.995	0.813	0.981	0.982	0.998
	neutral	0.000	0.000	0.004	0.001	0.000
	negative	0.005	0.188	0.015	0.017	0.002
	observations	918,940	64	476	206,040	712,360

The table contains sample probabilities for the type of first feedback along with sample frequencies for the conditioning set in columns (2) through (5) and in the second and third panel. In columns (2) through (5), we condition of experience patterns. We call a player experienced if the feedback score is at least 20. In the second and third panel, we additionally condition on the role of the first feedback giver.

Table 4: Types of first feedbacks and dependence on experience and rating sequence.

		reply is positive	neutral	negative	missing	observations
		0.697	0.003	0.005	0.295	2,470,522
feedback is	positive	0.709	0.002	0.002	0.288	2,428,335
	neutral	0.044	0.096	0.042	0.818	20,459
	negative	0.028	0.011	0.326	0.635	21,728
buyer rates first		0.702	0.001	0.005	0.291	1,551,582
and feedback is	positive	0.718	0.000	0.000	0.281	1,514,390
	neutral	0.044	0.097	0.042	0.817	20,103
	negative	0.029	0.013	0.395	0.563	17,089
seller rates first		0.690	0.005	0.004	0.301	918,940
and feedback is	positive	0.693	0.005	0.004	0.298	913,945
	neutral	0.031	0.045	0.053	0.871	356
	negative	0.023	0.005	0.071	0.901	4,639

The table contains sample probabilities for the type of second feedback (reply) conditional on the type of first feedback in the first panel. The first row in each panel contains probabilities for the second feedback irrespective of the type of first feedback. In the second and third panel, we condition on the role of the first feedback giver. The last column contains, by row, the respective number of observations.

Table 5: Second feedbacks as a function of first feedbacks and rating sequence.

3.3 Feedback Timing

Figure 2 shows the distribution of the time a feedback is left in the top row and the time difference between feedback and reply in the bottom row. It is interesting to compare instances in which a first feedback is left by the seller to the ones in which it is left by the buyer. In the former case, the rating is given much earlier, probably because of the usual transaction sequence on eBay, according to which payment precedes shipping. This is also reflected in the time difference between feedback and reply. Interestingly, the distribution for the latter is bimodal for buyers. If the first feedback is left by the seller, the buyer either responds relatively early with a reply, possibly reflecting the fact that she was waiting for the seller rating in order to leave a safe reply, or it takes some time until she gets the good and rates. Conversely, sellers often respond immediately to a first rating by a buyer, reflecting their desire not to expose themselves to the mercy of the buyer. In both cases, the timing of immediate replies can be interpreted as strategic. In fact, if he rates at all, the seller should be much more likely to rate first as opposed to replying, given the transaction sequence on eBay. This conflicts with the observation that the seller only rates first in 37 per cent of the cases. Interestingly, the shape of the time distribution is independent of the type of first feedback as we show in Figure 9, 10, and 11 in Appendix B. Moreover, the presence of immediate replies does not depend on the time a first feedback is left as we illustrate in Figure 3. While this pattern *per se* does not necessarily suggest bias in the construction of the reputation record, it clearly indicates that feedback giving by one party is dependent on when the other party is giving feedback, apparently independently of the flow of transactions. We conclude that this provides for potentially substantive bias also in the feedback category.

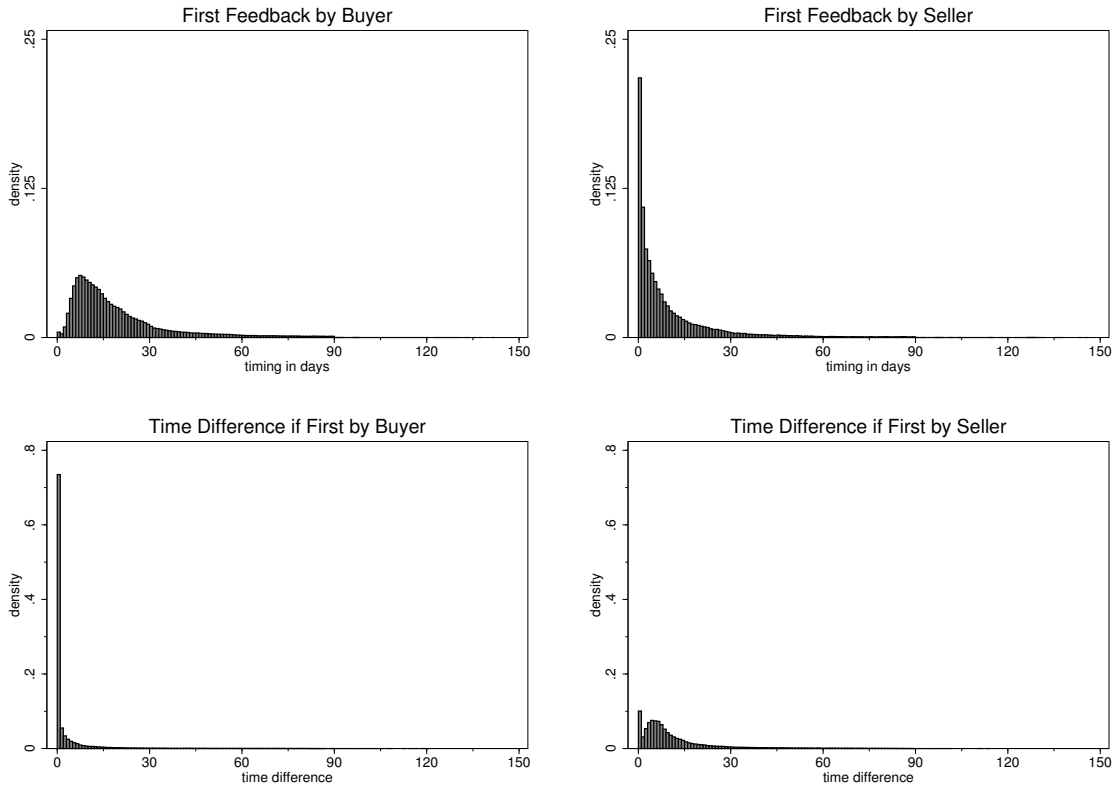

Figure 2: Feedback timing.

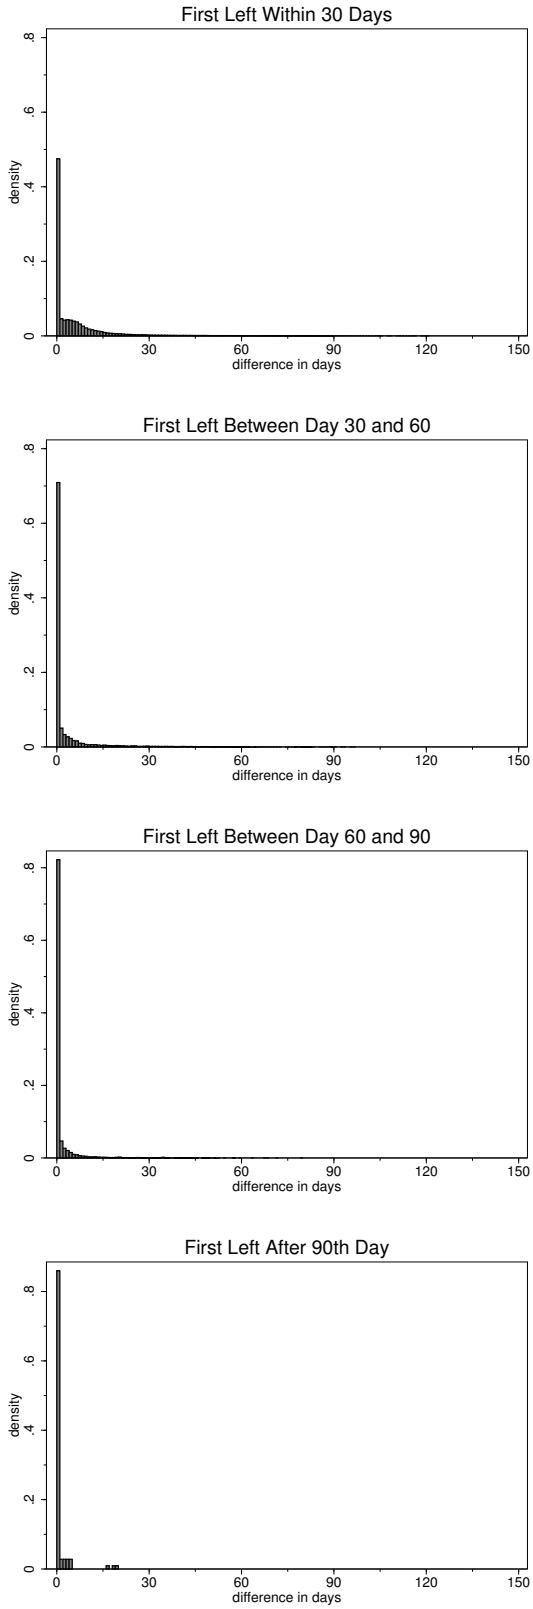

Figure 3: Timing difference between feedback and reply.

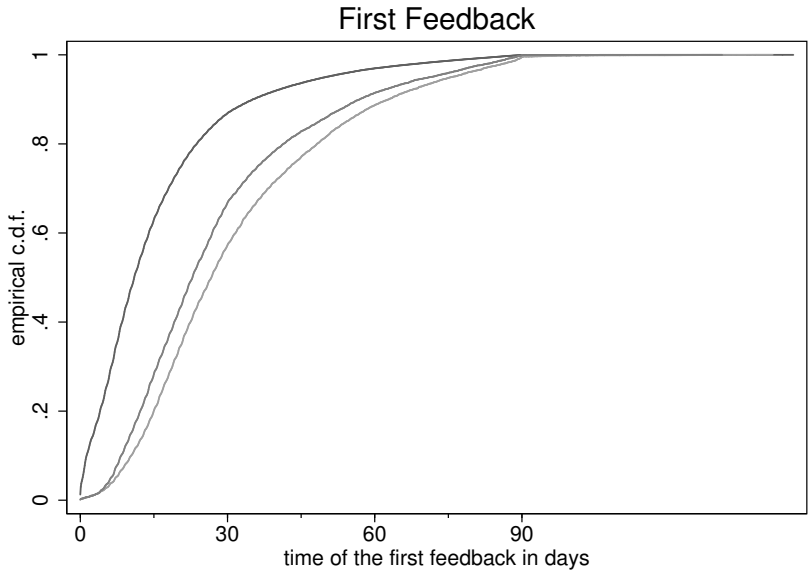

Figure 4: Empirical cumulative distribution functions for the timing of the feedback given that it is positive, neutral, or negative (from left to right). Note that negative and neutral feedback is given later in a first order stochastic dominance sense.

Whereas the shape of the distribution of the timing of first feedbacks does not depend on its type the average time that passes until a feedback is left does depend on its type. This is shown in Figure 4 which contains empirical distribution functions. Observe that in a first order stochastic dominance sense feedback is given earlier if it is positive rather than neutral, and in turn is given earlier if it is neutral rather than negative. With respective p-values of 0 this is confirmed by one-sided Kolmogorov-Smirnov tests.

However, even if feedback behavior was non strategic, negative or neutral marks could be given later simply because the transaction was delayed and *therefore*, a negative or neutral feedback was left. Hence, the documented effects could entirely stem from transactions characterized by late delivery in which a truthful negative report is posted late. Conversely, those transactions characterized by timely delivery on both sides are likely to produce truthful positive feedbacks that are posted early. Resnick and Zeckhauser (2002, Table 2) have analyzed the feedback comments belonging to a sample of negative or neutral marks. On one hand, they find that 11 per cent of the complaints were about slow shipment. Additionally, in 23 per cent of the cases buyers claimed not to have received the item after they had paid for it. Hence, there is at least some scope for delays. On the other hand, however, in 24 per cent of the cases the good was shipped in time but was in poor condition, thus giving room for truthful negative and timely feedback.

Figure 5 and 6 differentiate by the role and experience of the first feedback giver. They show that the time gap between first positives and negatives becomes wider for sellers and

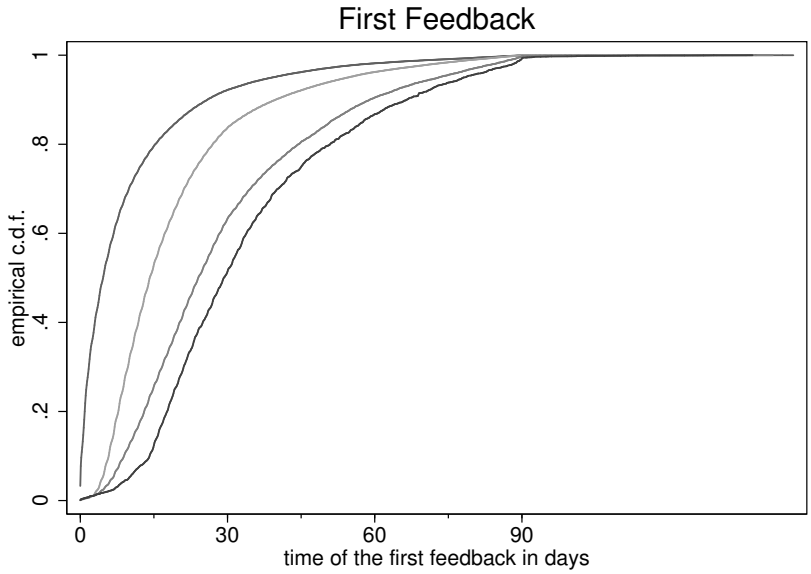

Figure 5: Empirical cumulative distribution functions for the timing of the feedback given that it is given by the seller and that it is positive, buyer and positive, buyer and negative or neutral, seller and negative or neutral (from left to right).

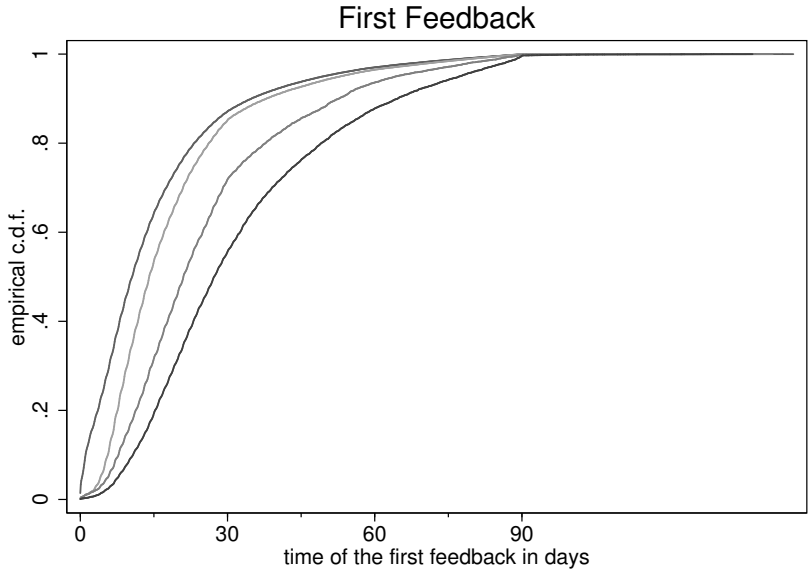

Figure 6: Empirical cumulative distribution functions for the timing of the feedback given high experience and the feedback being positive, low experience and positive, low experience and negative or neutral, high experience and negative or neutral (from left to right). High experience is defined as a feedback score of at least 20.

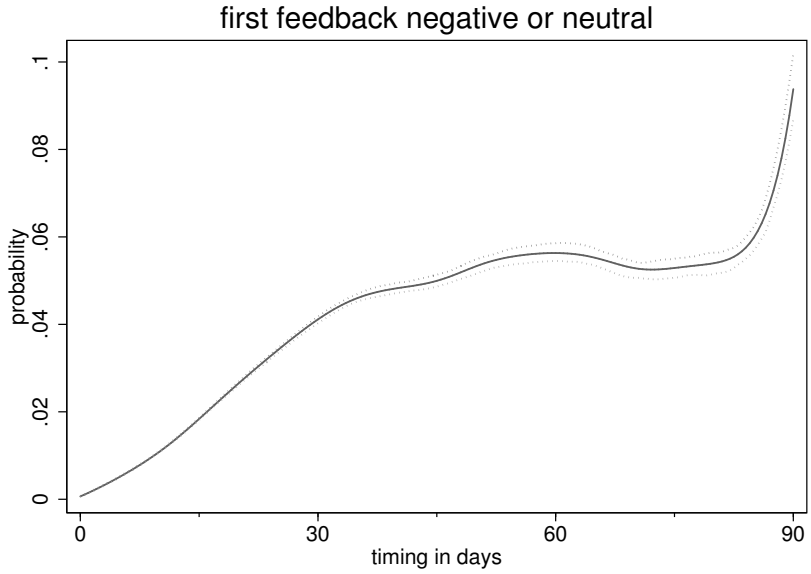

Figure 7: The probability of a neutral or negative feedback against time. Local linear regressions and bootstrapped 95 per cent confidence intervals (100 replications).

more experienced players.[32] One explanation could be that both sellers and more experienced players are more concerned about their reputation and thus have an incentive to act strategically by posting positive feedbacks early in order to induce reciprocation and negative ratings late in order to minimize the probability of receiving a retaliative negative feedback in return. A strategic consideration in this context could be that the attention of the trading partner is fading in time so that a negative feedback is less likely to be retaliated if it is given late.

3.4 Last Minute Feedback

We have argued that negative first feedbacks might be given late because it takes time until a user realizes that the transaction did not go well. Figure 7 shows the estimated conditional probabilities of the feedback being neutral or negative conditional on the time the feedback is left. This estimate was obtained using nonparametric local linear regressions of an indicator variable for negative or neutral feedback on the time it was left.[33] Figure 7 shows that the later the feedback is left, the more likely it is to be negative or neutral—even culminating into a spike around the 90th day which is apparently perceived by some users to be the last day feedback can be left. Arguably, users know whether a transaction went well after about 30-45 days as the probability of a negative or neutral mark seems not to depend on the timing of the feedback between then and the 90th day of the feedback period. However, it increases close

[32]This is confirmed by one-sided Kolmogorov-Smirnov tests with respective p-values of 0.

[33]We used a Gaussian kernel. It turned out that the choice of the bandwidth did not have a substantial impact on these estimates. Here, we chose them *ad hoc*. Notice that the bootstrapped confidence intervals are extremely narrow due to the size of the data set.

to the 90th day. We interpret this as evidence for strategic behavior by users who believe in a deterministic last minute of the feedback period. As for statistical inference, regressions show that this increase is highly significant at any level. This can hardly be reconciled with non strategic behavior since that would require that all of a sudden more negative or neutral than positive information on the trading partner would be revealed on the second half of the 90th day, compared to the 60 day period preceding this day.

Next, we evaluate in more detail how the institutional rules are perceived by users on eBay and how successful it is to leave a last minute feedback. First, notice that Figure 4 and the bottom row in Figure 9 show that most negative feedback is left relatively early within the 90 day period. Moreover, the latter figure shows that the spike in the frequency of feedback being left on day 90 as compared to the preceding days is relatively small. We therefore conclude that not many users leave feedback in the last minute. However, if they do so, it is much more likely to be negative.

In Table 6 we compare characteristics of last minute feedback instances to the ones of other instances. Column (1) and (2) show that last minute feedbacks are much more likely to be negative and given by the buyer. As compared to all types of feedbacks, last minute negative feedbacks are more likely to be given by experienced players as compared to all other negative feedbacks.

In Figure 8, we investigate whether leaving a negative feedback late is an empirically successful strategy to minimize the likelihood of receiving a negative feedback in return. The plots show the dependence between the reply and both the time of the feedback and its type. All graphs show that the later the feedback is given the less likely it is that a reply is given *at all*. More precisely, the probability that a reply is missing is increasing in time. This observation is independent of the type of the feedback. In passing, this also shows that early positives are more likely to be reciprocated. Hence, leaving a positive feedback early appears as an empirically successful strategy in terms of maximizing the probability of receiving a reciprocative positive in return.

Last, we investigate whether leaving a negative feedback on the last day of the feedback period decreases the likelihood of receiving a negative feedback in return. Table 7 shows probit marginal effects for the probability that a negative feedback is retaliated. It shows that buyers retaliate less often, that the probability to observe a retaliative negative increases in the experience of the second feedback giver, that it decreases in the time a first feedback is left, but is statistically unchanged if it is left in the last minute. This means that the institutional rules on eBay do in fact not allow users to leave a safe negative at any point in time, in line with the discussion of the institutional details in the Introduction and Section 2.

In this section, we have presented empirical evidence on the sequencing and timing of feedbacks that documents strategic rating behavior on eBay which potentially distorts feedback records. In the next section, in light of these findings, we try to quantify these biases.

4 Bounds

As transaction outcomes are not observed on eBay it is challenging to point identify the joint distribution of buyer and seller satisfaction. In this section, we derive Manski (2003) type

	(1)	(2)	(3)	(4)
			first negative	first negative
	not LMF	LMF	not LMF	LMF
first feedback positive	0.983	0.904	0.000	0.000
first feedback neutral	0.008	0.032	0.000	0.000
first feedback negative	0.009	0.064	1.000	1.000
first left by seller	0.372	0.244	0.213	0.321
feedback score of player giving first feedback	11,334.318	4,884.889	2,318.637	4,690.664
feedback score of player giving second feedback	16,859.573	20,858.533	27,377.805	23,880.575
first feedback giver experienced	0.877	0.858	0.643	0.873
second feedback giver experienced	0.916	0.920	0.837	0.761
both experienced	0.792	0.778	0.481	0.634
none experienced	0.000	0.000	0.001	0.000
observations	2,447,647	2,089	21,394	134

Note: The number of observations is the number of observations for which all variables are non-missing.

Table 6: Relationship between characteristics and LMF instances.

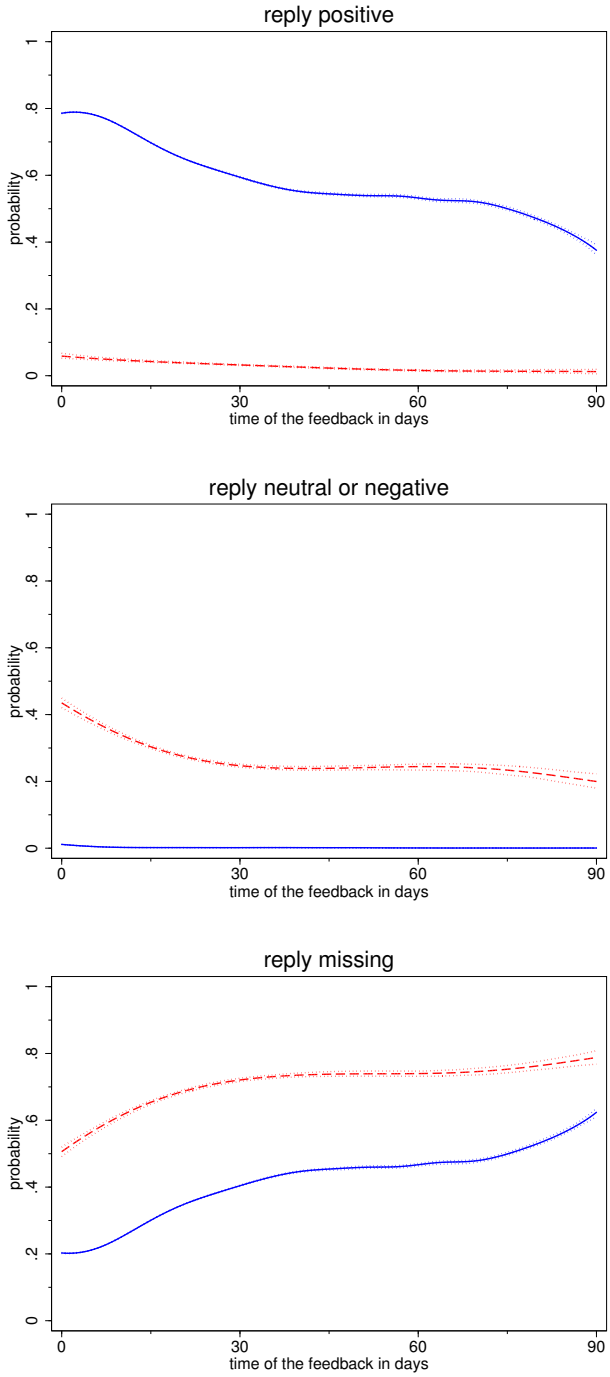

Figure 8: The probability of a positive (top), negative or neutral (middle), and missing (bottom) reply given a positive (solid line) and negative or neutral (dashed line) feedback against time of the feedback. Local linear regressions and bootstrapped 95 per cent confidence intervals (100 replications).

	negative retaliated
first feedback left by seller	-0.177**
	(0.012)
receiver of first feedback experienced	0.264**
	(0.011)
timing of first feedback	-0.003**
	(0.000)
timing of first feedback times indicator if within first 30 days	-0.003**
	(0.000)
last minute feedback	-0.038
	(0.045)
observations	21,728

Marginal Effects. Standard errors in parentheses. * significant at 5%; ** significant at 1%.

Table 7: Probit model for retaliation of a first negative feedback.

bounds for the satisfaction of buyers and sellers. One advantage of deriving bounds is that the assumptions we need to rely on are considerably weaker than the ones that are needed for point identification.

We exploit that rating behavior is related to the outcome of the transaction. However, at the same time, it is potentially led by strategic considerations. For example, a buyer may consider leaving a positive mark even if she is unsatisfied with the transaction once she believes that the seller will reciprocate and leave him a positive in return. This belief might be justified if the buyer himself performed well. Another example are retaliative second negatives which are given in order to be able to renegotiate for a withdrawn feedback, not because the buyer is unsatisfied with the transaction *per se*. Importantly, such strategic considerations are ruled out in the approach taken by Dellarocas and Wood (2006). In particular, they only allow for strategically delaying or not leaving a feedback.

Denote the satisfaction of buyer or seller $i \in \{b, s\}$ by $u^i \in \{h, m, l\}$, where h, m, and l denote high, medium, and low satisfaction, and his rating by $r^i \in \{+, 0, -\}$, where $+$, 0, and $-$ denotes a positive, neutral, or negative mark, respectively. Moreover, let h correspond to $+$ under truthful reporting, 0 correspond to m, and $-$ correspond to l. This implies that a truthful rating does not depend on the rating of the trading partner. Put differently, truthful reporting is defined as reporting behavior that is solely driven by the private knowledge on the outcome of the transaction. Finally, the rating sequence is denoted by the tuple $\sigma \in \{(b, s), (s, b), (s, \cdot), (b, \cdot), (\cdot, \cdot)\}$, where a $\cdot$ represents a missing rating so that, e.g. $(s, \cdot)$ denotes the case in which only the seller rates.

We now derive bounds for buyer satisfaction under different sets of assumptions. The derivation of the bounds for seller satisfaction is analogous. By the law of total probability, we can express the probability that the buyer is satisfied with the outcome of the transaction as a func-

tion of the rating sequence as

$$\Pr(u^b = h) = \sum_{\sigma \in \{(b,s),(b,\cdot)\}} \Pr(\sigma) \cdot \Pr(u^b = h|\sigma) + \sum_{\sigma \in \{(s,b),(s,\cdot)\}} \Pr(\sigma) \sum_{r^s \in \{+,0,-\}} \Pr(u^b = h|\sigma, r^s) + \Pr(\sigma = (\cdot,\cdot)) \cdot \Pr(u^b = h|\sigma = (\cdot,\cdot)). \quad (1)$$

That is, the probability that the buyer is happy is the sum of the probability that the buyer is happy conditional on him rating first, the probability that she is happy if she rates second for all possible first ratings by the seller, and the probability that the buyer is happy if no rating is observed at all.

Note that the transaction sequence is observable so that $\Pr(\sigma)$ is observable for all σ. However, as our data set contains only auctions for which we observe at least one feedback, we have to make an assumption on the probability that no feedback is left at all. Previous empirical analyses have shown that at least one feedback is left in 70 (Resnick and Zeckhauser, 2002), 60 (Chwelos and Dhar, 2005), and 88 (Dellarocas and Wood, 2006) per cent of the cases. Here, we report results that are based on the assumption that feedback is left in 70 per cent of the cases.

Next, we motivate several sets of assumptions on the truthfulness of positive buyer marks, depending on the rating pattern that is observed, and fully exploit them in order to bound the probability of buyer satisfaction. Bounds are calculated using the sample frequencies and sample probabilities that are reported in Table 9.

Scenario 1: All Positives are Truthful

The first set of bounds is derived under the assumption that ratings are either left truthfully or not at all. This assumption is very restrictive in that it assumes that only the decision whether or not to leave a rating is strategic, but not the decision about the type of rating. This excludes that cases that were discussed above, namely opportunistic first positives and retaliative second neutral or negative marks. However, our assumption is weaker than the assumptions made by Dellarocas and Wood (2006) who in addition make assumptions on the distribution of feedback timing under which the distribution of buyer satisfaction is point identified.

If positives can be assumed to be truthful, $\Pr(u^b = h|\sigma = (b,s) \vee \sigma = (b,\cdot)) = \Pr(r^b = +|\sigma = (b,s) \vee \sigma = (b,\cdot))$, where the right hand side (RHS) is observed in our data. Similarly, $\Pr(u^b = h|(s,b)) = \Pr(r^b = +|\sigma = (s,b))$.

We obtain the bounds for buyer satisfaction from the fact that the remaining unknown probabilities in (1), $\Pr(u^b = h|\sigma = (s,\cdot))$ $\Pr(u^b = h|\sigma = (\cdot,\cdot))$ lie between 0 and 1. The resulting bounds for buyer satisfaction are 60.9 and 98.7 per cent. As our data set consists of $2,470,522$ observations these bounds are estimated very precisely with a standard deviation of 0.03 and 0.01 per cent, respectively. Therefore, in the remainder, we abstain from reporting confidence intervals.[34]

[34]The bounds can be interpreted as Bernoulli random variables. If the bound is θ, the variance of the estimator for θ is given by $\theta(1-\theta)/N$ and can be estimated by the sample analogue.

	buyer satisfaction		seller satisfaction	
	lower bound	upper bound	lower bound	upper bound
Scenario 1	60.9%	98.7%	56.8%	99.6%
Scenario 2	18.0%	98.7%	30.9%	99.9%
Scenario 3	42.0%	92.7%	54.9%	93.9%
percentage positive marks if left	97.9%		99.3%	

Table 8: Bounds for buyer and seller satisfaction and percentage positive feedbacks left.

Scenario 2: Worst Case Bounds

Our second scenario is meant to illustrate the potential effects of strategic behavior. The underlying assumption is that only second positive feedbacks are truthful, i.e. $\Pr(u^b = h|\sigma = (s,b)) = \Pr(r^b = +|\sigma = (s,b))$, since first positives could be given strategically and second neutral and negative marks could be retaliative if given after a first neutral or negative mark. Under this assumption the bounds for buyer satisfaction are given by 18.0 and 98.7 per cent. Notably, these bounds are sharp in the sense that they constitute all that can be learned from the data without making any further assumption.

Scenario 3: Fractional Dissatisfaction

For the third set of bounds, in addition to the assumptions made in Scenario 2, we assume that in instances in which no feedback is left by both trading partners, buyers are happy with the transaction in 80 per cent of the cases so that $\Pr(u^b = h|\sigma = (\cdot,\cdot)) = 0.8$. This assumption is based on summary statistics reported in Chwelos and Dhar (2005). They compare feedbacks that are received by a commercial seller who sells the same products on Amazon and eBay. They report that the shop receives 13.5 per cent negatives on Amazon and 9.4 per cent on eBay. Combining these numbers with the fact that feedback is provided by buyers in 60.1 per cent of the cases on eBay and assuming that reports on Amazon are unbiased implies that 19.68 per cent of the missing feedbacks would have been negative. Under the assumption that it is equal to 20 per cent, our bounds indicate that the probability that a buyer is satisfied with a transaction lies between 42.0 and 92.7 per cent.

Discussion

Table 8 contains the derived sets of bounds for buyer satisfaction, along with bounds for seller satisfaction, which were calculated analogously, as well as the percentage of positive marks among the ones that were given. The worst case bounds that were derived in Scenario 2 indicate that in the presence of strategic non- and misreporting it highly depends on the beliefs about marks that are missing whether the observed percentage of positive marks on eBay directly reflects buyer and seller satisfaction. Notably, the upper bound both for buyer and seller satisfaction under Scenario 3 lies considerably below the percentage of positive marks observed on eBay. Under the assumptions made, this is evidence for substantial bias in eBay feedback scores.

5 Policy Implications

We have argued in this paper that not only feedback timing, and hence the rating sequence, but also the type of feedback that is left might be led by strategic considerations. Given the mechanism implemented by eBay, that allows for the observation of feedbacks given by the trading partner at any time, positive first feedback may be given strategically but untruthfully, because it removes potential hold up on the trading partner and thus induces him to reciprocate. By contrast, leaving a truthful negative feedback is a risky endeavor because it may be retaliated untruthfully, and purely for strategic reasons. While the existence of a deterministic last minute of the feedback period would allow for unretaliated negative first feedback giving, the end of the feedback period is stochastic. Moreover, our empirical analysis led us to discover a previously unnoticed feature of the feedback mechanism: Whenever a transaction is rated, the time window during which a reply can be left for this transaction is in fact extended.

In consequence to this, let us develop some ideas towards improving on the design of the feedback mechanism.[35] Our analysis suggests that to reduce concerns for retaliation and foster expression of deserved dissatisfaction the "feedback game" should be made less transparent to both parties. In particular, favorable "anonymity" should be pursued, so that both feedbacks are revealed to the trading partners and the public only if no more feedbacks can be left.[36] This could be done after a fixed period, or after both have already given their feedback. Note that this device requires that feedback withdrawal is not possible. Otherwise, under general conditions, it remains a dominant strategy for the players to always leave a negative feedback in order to be able to renegotiate after feedbacks have been revealed.

In general, the performance of buyers, if asked to pay first, is subject to little uncertainty. Either the full payment arrives in time, and bank transfer details can demonstrate this, or it does not. Thus, if the sequence of transactions is such that payments preceed delivery, opportunism on the buyer's side does not play a role. However, sellers can misbehave on a variety of aspects of their performance, and this opaqueness creates room for opportunistic behavior. Therefore, it may be worthwhile to limit feedbacks to buyers rating sellers as in Amazon auctions.[37] In equilibrium this should induce the sequencing of transactions suggested above.

While such a change of rules would most likely result in less (positive) feedback giving, it should (thereby) lead to removing substantive biases in the reputation records in this increasingly important market.

Appendix A: Data Collection

We first randomly drew auction numbers and downloaded the respective auction details. From these auction details we obtained the respective seller member ID and randomly selected 10,000 sellers from the United States.

In a next step, for each seller, we used the information in her feedback profile to obtain auction details including the corresponding feedback which was received and left, and the re-

[35] Roth (2002) makes a strong case for economists helping to *design* markets and institutions.

[36] This has also been suggested by Reichling (2004).

[37] This is also suitable for e-procurement platforms. See Dini and Spagnolo (2005a,b) for further details.

spective timing information. By construction, since we start from a member's feedback profile, our sample consists of auction records for which at least one feedback was left by either the seller or the buyer. In order to minimize the loss of information, we included only those auctions into our data set which ended at least 100 days before the date of our data collection.

We restricted our attention to standard eBay auctions. That is, we dropped listings that belong to "eBay Motors," are "Live Auctions," serve as an advertisement only (with no bidding possible) and are lot auctions (in eBay jargon: "Dutch" auctions). Moreover, we did not consider auctions that ended early.

Records containing mutually withdrawn feedbacks were dropped.

Appendix B: Additional Tables and Figures

E_1	$\Pr(E_1)$	N_1	E_2	$\Pr(E_2\|E_1)$	N_2	E_3	$\Pr(E_3\|E_1,E_2)$	N_3
first by buyer	0.63	1,551,582	buyer positive	0.98	1,514,390	seller positive	0.72	1,087,988
						seller neutral	0.00	118
						seller negative	0.00	173
						seller missing	0.28	426,111
			buyer neutral	0.01	20,103	seller positive	0.04	889
						seller neutral	0.10	1,949
						seller negative	0.04	847
						seller missing	0.82	16,418
			buyer negative	0.01	17,089	seller positive	0.03	494
						seller neutral	0.01	225
						seller negative	0.39	6,749
						seller missing	0.56	9,621
first by seller	0.37	918,940	seller positive	0.99	913,945	buyer positive	0.69	633,492
						buyer neutral	0.00	4,546
						buyer negative	0.00	3,596
						buyer missing	0.30	272,311
			seller neutral	0.00	356	buyer positive	0.03	11
						buyer neutral	0.04	16
						buyer negative	0.05	19
						buyer missing	0.87	310
			seller negative	0.01	4,639	buyer positive	0.02	108
						buyer neutral	0.00	21
						buyer negative	0.07	330
						buyer missing	0.90	4,180
observations		2,470,522			2,470,522			2,470,522

Event E_1 denotes whether a first feedback is left by the buyer or seller. $\Pr(E_1)$ denotes the sample probability for the event E_1. N_1 is the corresponding number of observations. Event E_2 denotes the type of first feedback and $\Pr(E_2|E_1)$ denotes the sample conditional probability for E_2 conditional on E_1. N_2 is the corresponding frequency of the joint event E_1 and E_2. E_3, finally, denotes the reaction by the trading partner, including the alternative not to rate. The next column contains the sample probability and the last column contains the frequency of the joint event E_1, E_2, E_3. Note that the table is conditional on at least one feedback being left.

Table 9: Feedback patterns.

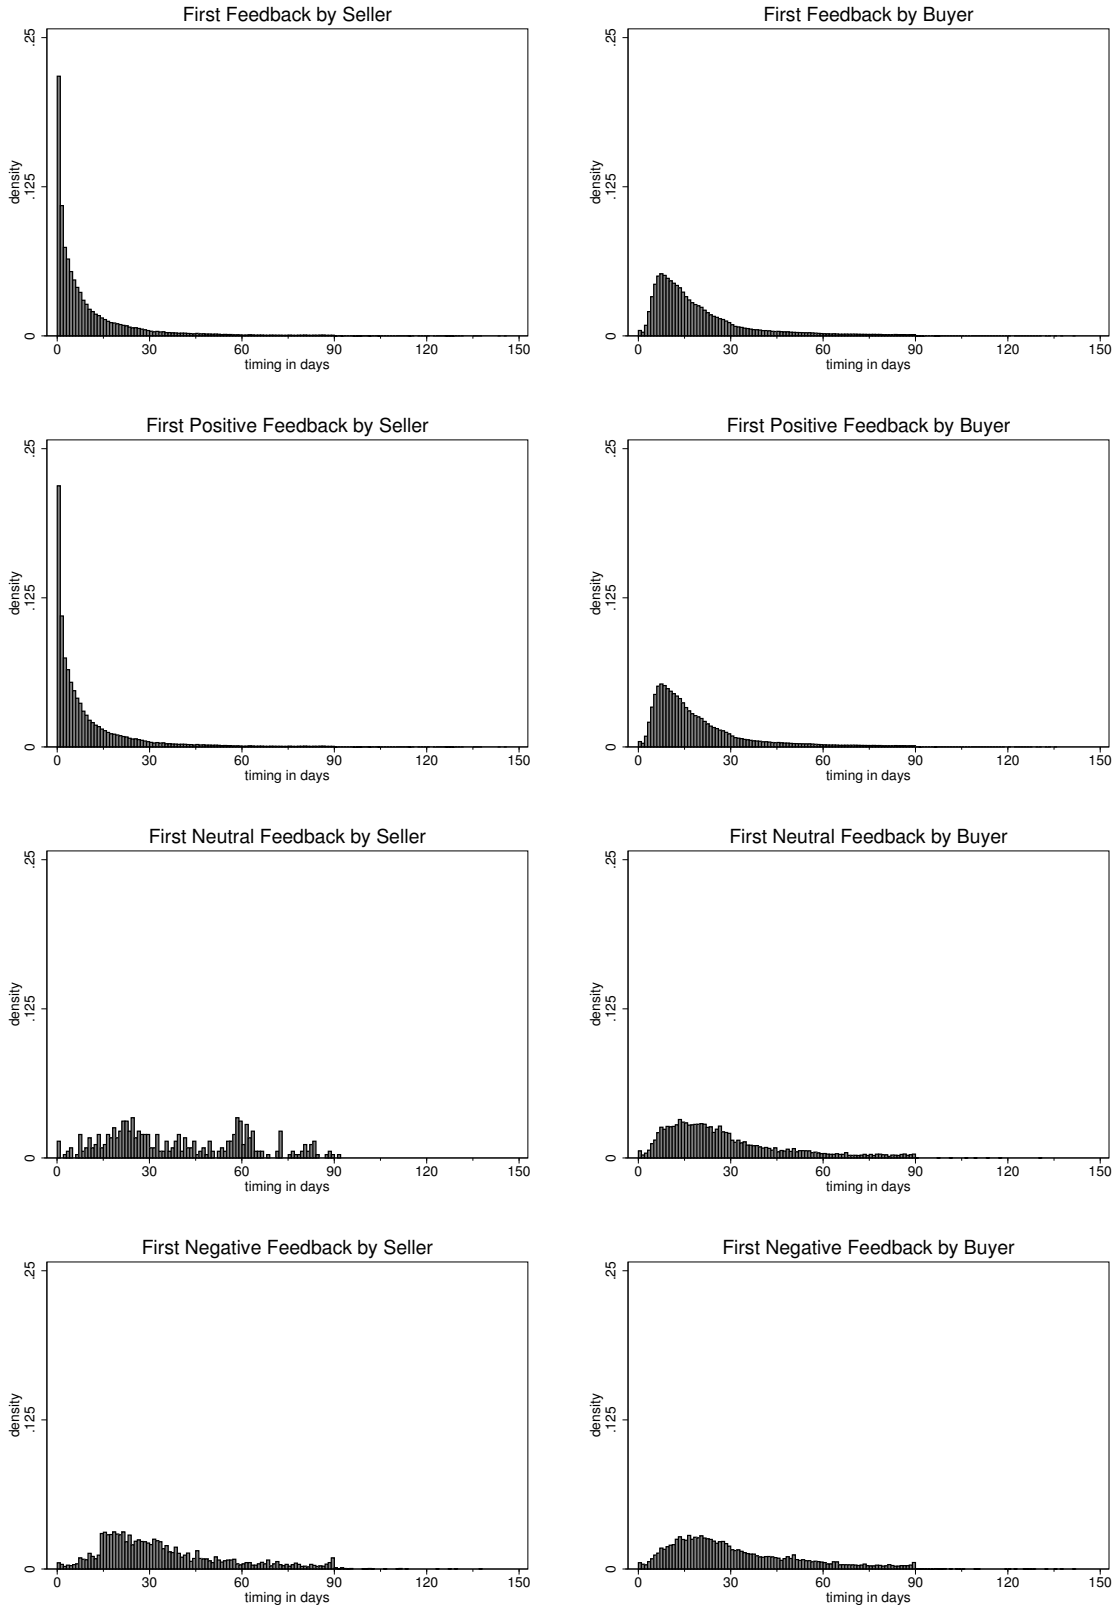

Figure 9: Timing of first feedbacks.

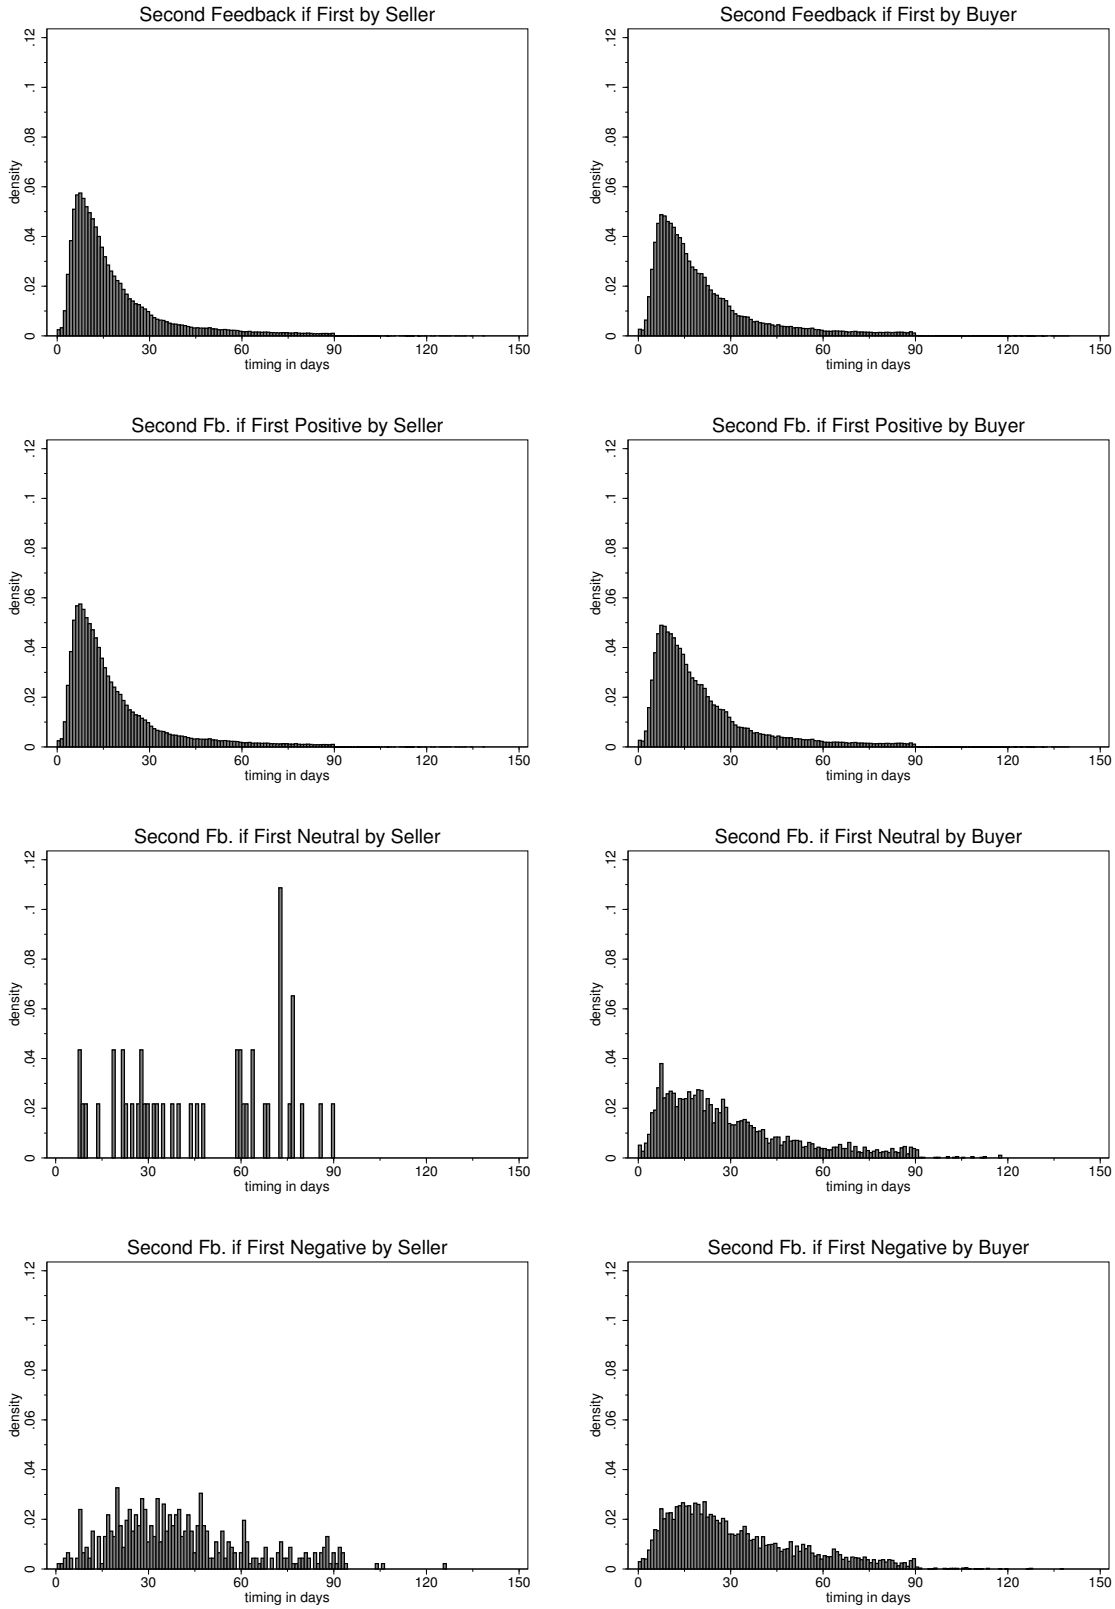

Figure 10: Timing of replies.

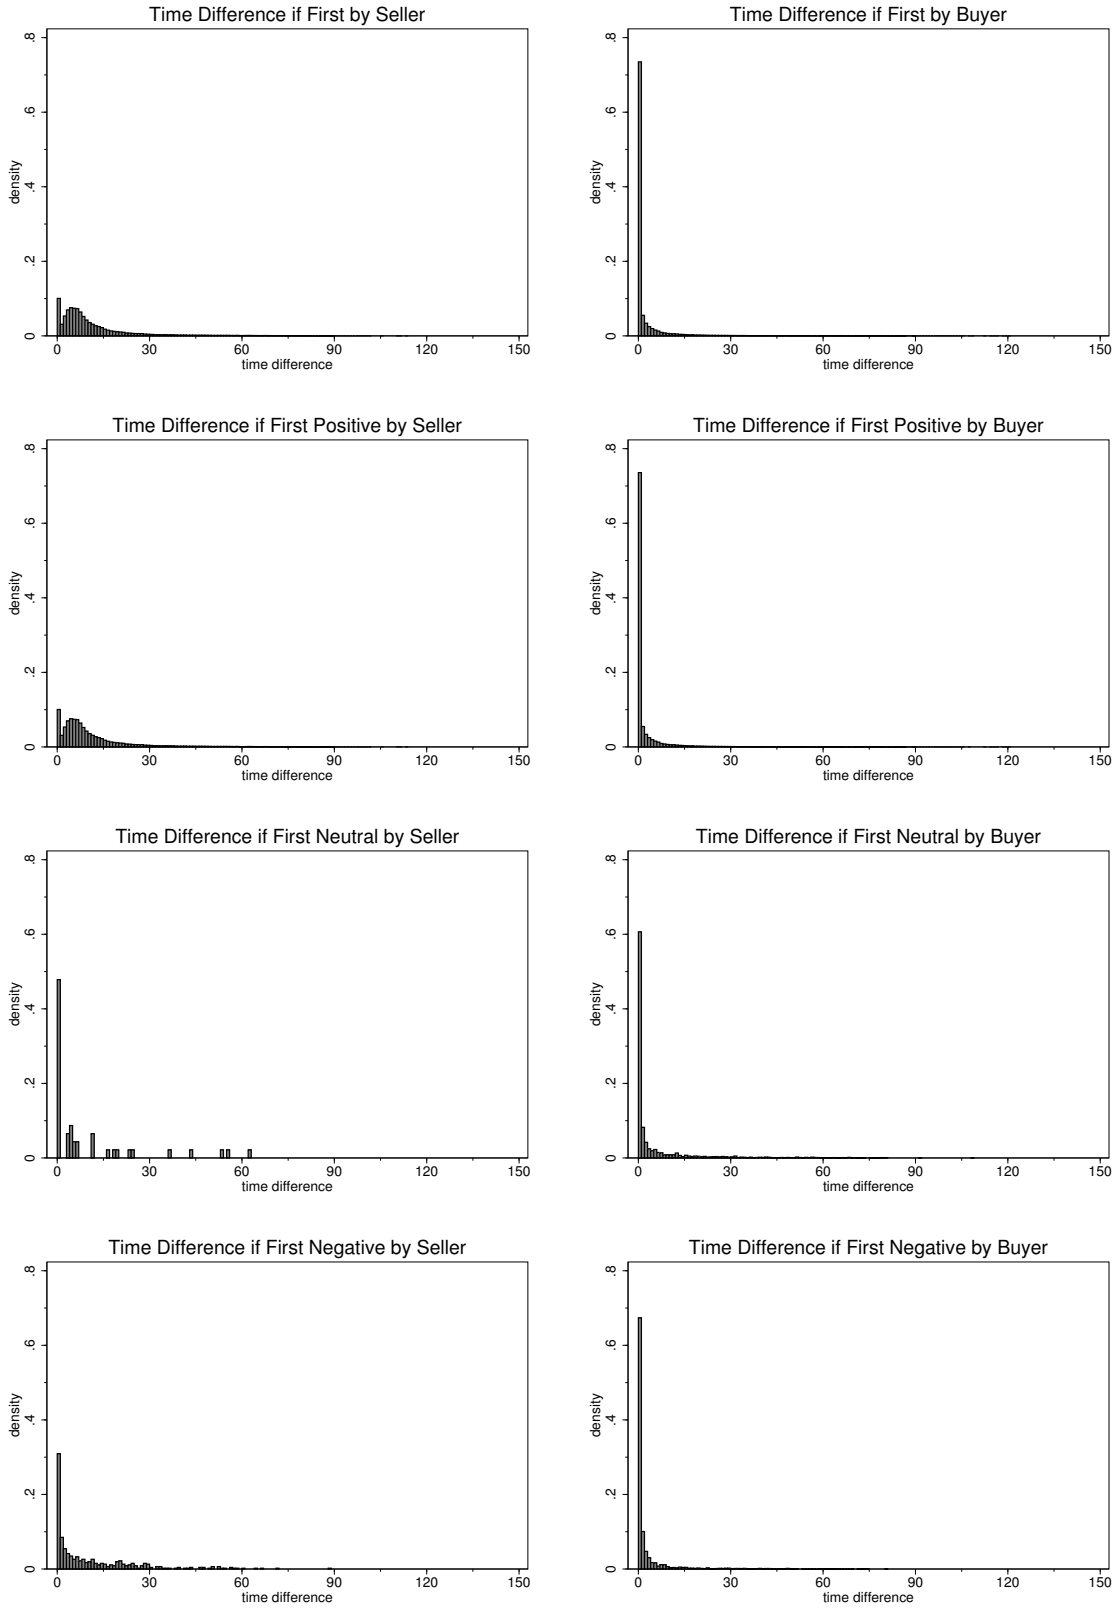

Figure 11: Timing difference between feedback and reply by type of first feedback and rating sequence.

References

BAJARI, P., AND A. HORTAÇSU (2003): "The Winner's Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions," *RAND Journal of Economics*, 34(2), 329–355.

——— (2004): "Economic Insights from Internet Auctions," *Journal of Economic Literature*, 42(2), 457–489.

CABRAL, L. M. B., AND A. HORTAÇSU (2006): "The Dynamics of Seller Reputation: Theory and Evidence from eBay," Mimeograph.

CHWELOS, P., AND T. DHAR (2005): "Caveat Emptor: Differences in Online Reputation Mechanisms," Mimeograph.

DELLAROCAS, C. (2005): "Reputation Mechanisms," Working Paper, University of Maryland.

DELLAROCAS, C., AND C. A. WOOD (2006): "The Sound of Silence in Online Feedback: Estimating Trading Risks in the Presence of Reporting Bias," Mimeograph.

DINI, F., AND G. SPAGNOLO (2005a): "Buying Reputation on eBay," Manuscript, Consip Research Unit and Stockholm School of Economics.

——— (2005b): "Reputation Mechanisms and Electronic Markets: Economic Issues and Proposals for Public Procurement," in *Challenges in Public Procurement: an International Perspective*, ed. by K. V. Thai, A. Araujo, R. Y. Carter, G. Callender, D. Drabkin, R. Grimm, K. R. Ejlskov Jensen, R. E. Lloyd, C. P. McCue, and J. Telgen. Academic Press.

FEHR, E., AND K. SCHMIDT (1999): "A Theory of Fairness, Competition and Cooperation," *Quarterly Journal of Economics*, 114(3), 817–868.

HOUSER, D., AND J. WOODERS (forthcoming): "Reputation in Auctions: Theory, and Evidence from eBay," *Journal of Economics and Management Strategy*.

JIN, G. Z., AND A. KATO (2002): "Blind Trust Online: Experimental Evidence from Baseball Cards," Working Paper, University of Maryland.

LIVINGSTON, J. A., AND W. N. EVANS (2004): "Do Bidders in Internet Auctions Trust Sellers? A Structural Model of Bidder Behavior on eBay," Working Paper, Bentley College.

LUCKING-REILEY, D., D. BRYAN, N. PRASAD, AND D. REEVES (2005): "Pennies from eBay: the Determinants of Price in Online Auctions," Working Paper, University of Arizona.

MANSKI, C. F. (2003): *Partial Identification of Probability Distributions*. Springer, New York.

MCDONALD, C. G., AND V. C. SLAWSON (2002): "Reputation in an Internet Auction Market," *Economic Inquiry*, 40(3), 633–650.

MELNIK, M. I., AND J. ALM (2002): "Does a Seller's Reputation Matter? Evidence from eBay Auctions," *Journal of Industrial Economics*, 50(3), 337–349.

REICHLING, F. (2004): *Effects of Reputation Mechanisms on Fraud Prevention in eBay Auctions*Mimeograph.

RESNICK, P., K. KUWABARA, R. ZECKHAUSER, AND E. FRIEDMAN (2000): "Reputation systems," *Communications of the ACM*, 43(12), 45–48.

RESNICK, P., AND R. ZECKHAUSER (2002): "Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay's Reputation System. The Economics of the Internet and E-Commerce.," in *Advances in Applied Microeconomics*, ed. by M. R. Baye, vol. 11, Amsterdam. Elsevier Science.

RESNICK, P., R. ZECKHAUSER, J. SWANSON, AND K. LOCKWOOD (2004): "The Value of Reputation

on eBay: A Controlled Experiment," Working Paper, Harvard Kennedy School of Business.

ROTH, A. E. (2002): "The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics," *Econometrica*, 70(4), 1341–1378.

ROTH, A. E., AND A. OCKENFELS (2002): "Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon on the Internet," *American Economic Review*, 92(4), 1093–1103.