

An Experimental Analysis of Ending Rules in Internet Auctions

Dan Ariely^{*}
MIT, USA

Axel Ockenfels
Max Planck Institute for Research into Economic Systems, Germany

Alvin E. Roth
Harvard University, USA

December 3, 2002

Abstract

A great deal of late bidding has been observed on internet auctions such as eBay, which employ a second price auction with a fixed deadline. Much less late bidding has been observed on internet auctions such as those run by Amazon, which employ similar auction rules, but use an ending rule that automatically extends the auction if necessary after the scheduled close until ten minutes have passed without a bid. This paper reports an experiment that allows us to examine the effect of the different ending rules under controlled conditions, without the other differences between internet auction houses that prevent unambiguous interpretation of the field data. We find that the difference in auction ending rules is sufficient by itself to produce the differences in late bidding observed in the field data. The experimental data also allow us to examine how individuals bid in relation to their private values, and how this behavior is shaped by the different opportunities for learning provided in the auction conditions.

* Ariely: Sloan School of Management, MIT, Cambridge, MA, USA; ariely@mit.edu (web.mit.edu/ariely/www/); Ockenfels: Max Planck Institute for Research into Economic Systems, Strategic Interaction Unit, Kahlaische Straße 10, D-07745 Jena, Germany; ockenfels@mpiew-jena.mpg.de (www.mpiew-jena.mpg.de/esi/ockenfels/); Roth: Harvard University, Department of Economics and Harvard Business School, 183 Baker Library, Boston, MA 02163, USA; aroth@hbs.edu (www.economics.harvard.edu/~aroth/alroth.html).

We thank Ernan Haruvy for his support in programming the software and Muriel Niederle for her helpful comments. Ockenfels gratefully acknowledges the support of the Deutsche Forschungsgemeinschaft through the Emmy Noether-program. Roth gratefully acknowledges the support of the National Science Foundation and of the Harvard Business School.

I. Introduction

How to end an auction is a subject of active concern in the auction design literature. The concern is that some rules give bidders an incentive to bid late, which hampers price discovery and efficiency.¹ Internet auctions provide new opportunities to examine the effects of ending rules on bidding behavior, because some of the internet auction houses such as eBay and Amazon use essentially identical rules except for the rule that governs auctions' endings. Of course there are other differences between eBay and Amazon than their rules, so the field data are not always simple to interpret. Here we present the results of a laboratory experiment designed to complement the field data by controlling away all differences except the difference in auction rules, in order to allow their contribution to the differences in bidding behavior to be unambiguously observed.

The difference between the eBay and Amazon auction rules is that:

- eBay auctions have a fixed deadline (a “hard close”), that is, they end at a scheduled time, most often after seven days.
- Amazon auctions, in contrast, are automatically extended if necessary, past the scheduled end time, until ten minutes have passed without any bid being submitted.

Roth and Ockenfels (2002) compare the timing of bids in eBay and Amazon second-price auctions on the internet, observing a marked difference in bidding behavior:²

- Bids in eBay auctions are much more concentrated near the end of the auction than bids on Amazon auctions.
- More experienced bidders (as measured by their “feedback ratings”) are more likely to bid late on eBay, and less likely to bid late on Amazon.

For example, more than two-thirds of the eBay auctions in their sample had bids submitted less than an hour before the scheduled end time, in contrast to less than one quarter of the Amazon auctions. In the last 10 minutes, only 11 percent of the Amazon

¹ For example, the FCC auctions of radio spectrum licenses include “activity rules” intended to prevent bidders from concentrating their serious bids only near the end of the auction (see e.g. Milgrom, 2002, Roth, 2002).

² After the Roth and Ockenfels (2002) data were collected, eBay changed some rules, so that a contemporary comparison of field data now involves some additional differences. For example, since October 2000, eBay's bid history for each auction includes all bids, but Amazon's data continue to include only each bidder's last bid. Also, eBay sellers now have the opportunity to modify the reservation price during the auction under certain circumstances.

auctions received bids (i.e. only 11 percent of the Amazon auctions were extended past the scheduled deadline), while more than half the eBay auctions received bids in the last ten minutes (and over 10% of the eBay auctions received bids in the last ten seconds).

Ockenfels and Roth (2001) demonstrate that there are multiple reasons why the difference in ending rules between eBay and Amazon might produce this difference in bidding behavior. They model a second price auction conducted over time in which early bids give other bidders time to respond, but can be submitted with certainty, while very late bids do not give other bidders time to respond, but have a chance that they will not be successfully transmitted.³ Ockenfels and Roth (2001) show that in such an environment late bidding (also known as “sniping”) can arise as a rational response to many different causes. In both private value and common value auctions, and both in equilibrium and as best response to incremental bidding, the ending rules create incentives to bid late on eBay, but not on Amazon. The observation that the bidding behavior between the two auctions differs in the predicted way lends support to the hypothesis that these strategic incentives induced by the auctions’ ending rules are the cause of the difference in bid timing.

Interpretation of such field data is complicated by the fact that there are differences between eBay and Amazon other than their ending rules. eBay has many more items for sale than Amazon, and many more bidders. Furthermore, buyers and sellers themselves decide in which auctions to participate, so there may be differences between the characteristics of sellers and buyers and among the objects that are offered for sale on eBay and Amazon. Some combination of these uncontrolled differences between eBay

³ Surveys of late bidders by Roth and Ockenfels (2002) identified at least two sources of risk involved in late bidding. One was that bidders who plan to bid late sometimes find that they are unavailable at the end of the auction. The other involves bidders who are attempting to bid at the last moment but who do not succeed due to, e.g., erratic Internet traffic or connection times. In particular, more than 80 percent of the bidders who successfully bid at least once in the last minute of an eBay auction replied that it happened at least once to them that they started to make a bid, but the auction was closed before the bid was received. This risk also applies to artificial bidding agents. Esnipe.com, for instance, offers to automatically submit a predetermined bid a few seconds before the end of the eBay auction, but cannot guarantee that the bids are actually placed. In fact, esnipe.com reports on the basis of more than 4200 bids per day, that on average 4.5 percent of esnipe's bids failed to be successfully transmitted in September 2000 (<http://www.esnipe.com/stats.asp>, 2000).

and Amazon might in fact be the cause of the observed difference in bidding behavior, instead of the differences in rules.⁴

Moreover, the “feedback ratings” used by Ockenfels and Roth (2001) as proxies for experience are imperfect.⁵ For example, feedback ratings only reflect the number of completed transactions, but not auctions in which the bidder was not the high bidder. In addition, more experienced buyers on eBay may not only have more experience with the strategic aspects of the auction, they may have other differences from new bidders, e.g., they may also have more expertise concerning the goods for sale.

While the field data suggest that strategic incentives cause late bidding on eBay, the data do not easily allow us to focus on how each of the multiple reasons for late bidding contribute to the differences in bidding behavior on eBay and Amazon. For instance, the fact that bids on eBay antique auctions are even more skewed towards the deadline than those in auctions of computers suggests that the information conveyed by bids may play a role in promoting late bids on eBay auctions for goods with common values.⁶ In auctions with common values, late bidding could result because bidders might change their own evaluation as a reaction to the information in others’ bids. Similarly, bidders might want to bid late in order not to convey their information to others. However, the fact that the difference between eBay and Amazon auctions is clear even for auctions of computers seems to suggest that the different ending rules elicit different strategic incentives also in private value auctions. Our experiment will test the theoretical prediction that the hard close creates incentives to bid late even in the case of purely private values.

We report laboratory experiments on second-price auctions that differ only in the rule for how the auctions ended. Subjects were randomly assigned to each auction type, so there were no systematic differences in bidder characteristics across auctions, and the number of bidders per auction was kept constant. Each bidder in the experiment participated in a sequence of auctions, allowing us to observe in detail how bidding

⁴ Roth and Ockenfels (2002) briefly discuss possible self selection effects and their potential impact on bidding behaviors on eBay and Amazon. Another uncontrolled reason for late bidding in the field is that a bidder who is interested in an item available simultaneously in multiple auctions may want to delay his decision in which auction to bid.

⁵ Dellarocas (2002) presents a theoretical study, Resnick and Zeckhauser (2002) a field study, and Bolton et al. (2002) an experimental study that examine the problems and merits of eBay’s “feedback forum.”

⁶ See Bajari and Hortaçsu (2000) and Ockenfels and Roth (2002) who have explored the incentives that bidders have to bid late in common value auctions.

changes as bidders gain experience with the auction environment. The goods offered in our auctions were artificial, independent private-value commodities (each bidder was given a redemption value he would be paid in cash if he won the auction, and these values were drawn independently of the values of other bidders).

We experimentally compare several kinds of auctions, to allow us to separate different explanations of late bidding. By using our experimental setup we can also compare the revenues resulting from the different types of auctions (which could not be meaningfully compared in field data involving different commodities and different numbers of bidders), and the relative efficiency of the different auctions (since in the lab we will know in each auction which bidder has the highest value for the item auctioned).⁷

II. Experimental environment

II.1 The auction games

The treatments include four auction types: sealed bid, Amazon, eBay.8, and eBay1 (the latter two treatments differ only in the probability that a “last minute” bid will be transmitted). There were exactly two competing bidders in each auction. Each bidder in each auction was assigned a private value independently drawn from a uniform distribution between \$6 and \$10. The winner of an auction received his private value minus the final price, and a loser received nothing for that auction. The final price was determined by the second price rule, that is, the bidder who submitted the highest bid won and paid (at most) a small increment (\$0.25) above the highest bid of the opponent,

⁷ This study is part of a broader effort to understand the role of timing in transactions. Others have noted deadline effects in internet auctions (cf. Bajari and Hortaçsu, 2000, Malhotra and Murnighan, 2000, Wilcox, 2000), and similar deadline effects have been noted in studies of bargaining (cf. Gächter and Riedl, 2002, Güth et al., 2001, Roth et al., 1988, among others). Nonhuman animals have also been observed to change their behavior as a function of temporal distance from the end of the experience. Rats and pigeons respond more vigorously as the expected end of a fixed interval reinforcement schedule approaches (a pattern known as scalloping), even when the more vigorous behavior does not increase payoffs (Ferster and Skinner, 1957; Dews, 1969). Similarly, people who monitor the time of a task become increasingly impatient toward its end (Ceci and Bronfenbrenner, 1985), and people use different strategies when games are framed as getting close to the end (even when these are arbitrary break points; Croson, 1996). There are also environments in which the strategic use of time moves transactions earlier rather than later, and experiments have been conducted to better understand these incentives to transact early (see Haruvy et al., 2001, and Kagel and Roth, 2000, which are motivated by studies of timing in entry level labor markets, cf. Roth and Xing, 1994, and Avery et al. 2001).

or, if the opponent did not bid, the price was the minimum bid of \$1.⁸ All auctions were run in discrete time, so that we can precisely define ‘bidding late’ without running into problems of continuous time decision making such as individual differences in typing speed, which might differentially effect how late some bidders can bid.⁹

It will be easiest to describe the different auction conditions by first describing the eBay.8 treatment. It consists of two kinds of bidding stages, stage 1 (early) and stage 2 (late).

eBay.8

Stage 1 is divided into discrete periods. In each period, each trader has an opportunity to make a bid (simultaneously). At the end of each period, the high bidder and current price (minimum increment over second highest bid) are displayed to all. Stage 1 ends only after a period at which no player makes a bid.

Stage 2 of the eBay.8 auctions consists of a single period. The bidders have the opportunity to submit one last bid; it has probability $p = 0.8$ of being successfully transmitted.

eBay1

In the eBay1 condition, the probability that a bid made in stage 2 will be transmitted successfully is $p = 1$, i.e. stage 2-bids are transmitted with certainty. Everything else is as on eBay.8.

Amazon

⁸ The price never exceeds the highest submitted bid: If the difference between the highest and the second highest submitted bid is smaller than the minimum increment, the price paid is equal to the highest bid. If both bidders submitted the highest bid, the bidder who submitted his bid first is the high bidder at a price equal to the reservation price. If identical bids are submitted simultaneously, one bidder is randomly chosen to be the high bidder. Also, a bidder can bid against himself without penalty if he is the current high bidder, because it raises his proxy bid without raising the price.

⁹ Note that because eBay and Amazon are online auctions, it would have been possible to conduct the auction using precisely the eBay and Amazon interfaces, had that been desirable, by conducting an experiment in which the auctions were on the internet auction sites (for a classroom demonstration experiment of this sort, in a common value environment, see Asker et al. 2002). This would not have served our present purpose as well as the discrete version described here. In this respect it is worth noting that what makes an experimental design desirable is often what makes it different from some field environment, as well as what makes it similar.

Similar to the eBay.8 condition in the Amazon condition stage 1 is followed by stage 2 and the probability that a stage 2-bid will be successfully transmitted is $p = 0.8$. However, a successfully submitted stage 2-bid starts again stage 1-bidding (and is followed by stage 2 again, etc.). Thus, in the Amazon condition, the risk of bidding late is the same as in the eBay.8 condition, but a successful stage 2-bid causes the auction to be extended.

Sealed bid

In the sealed bid condition, the auction begins with stage 2 (with $p = 1$), and ends immediately after, so that each bidder has the opportunity to submit only a single bid, and must do so without knowing the bids of the other bidder. While the sealed bid auction obviously cannot yield any data on the timing of bids, it provides a benchmark against which behavior in different auctions can be assessed.

As in the internet counterparts, bidders in the eBay and Amazon conditions were always informed about current prices as the auction progressed, but the magnitude of the high bidder's current bid was never revealed to the low bidder.¹⁰ Also, each bid had to meet or exceed the current minimum acceptable bid, which was \$1 if no bid was submitted yet, or the smallest increment (\$0.25) over the current price or over one's own previously submitted bid (if any) whichever was higher.

Our experimental games reproduce the pricing and feedback policies employed by Amazon and eBay on the internet, and capture the essential differences in ending rules. First, there is sufficient time to submit bids and respond to others' bids early in the experimental conditions (that is, in stage 1). Second, there is a hard close in the eBay treatments that does not allow bidders to respond to very late (that is, stage 2) bids. The risk involved in submitting late bids in the eBay.8 condition reflects the fact that late bids run the risk of being lost in internet auctions (see Section I). eBay1, on the other hand, allows us to study the impact of this risk and therefore to separate different contributory causes of late bidding (as we will explain below). Third, successfully submitted late bids

¹⁰ In a situation in which the difference between the (current) price and the low bidder's (current) bid is smaller than the minimum increment, however, the low bidder can infer that the high bidder's (current) bid equals the price.

in the experimental Amazon condition automatically extend the auction (that is, move the auction back to stage 1), giving other bidders sufficient time to respond to all bids. However, late bidding on Amazon faces the same risk as late bidding on eBay.8. Finally, as in eBay and Amazon auctions on the internet, the second price rule allowed a bidder in the experiments to have his bid used to bid for him by proxy. That is, a bidder could submit a bid early in the auction and have the resulting price register as the minimum increment above the second highest bid. As the other bidder submits subsequent bids, the price rises to the minimum increment over the other player's bid until the bid is exceeded. Hence, as in the internet auction houses, an early bid that is higher than any other submitted during the auction will win the auction and pay only the minimum increment above the second highest submitted bid.

Table 2 summarizes our experimental auctions.

Auction condition	Number of stage 1 periods	Number of stage 2 periods	Probability of stage 2-bid to be successfully transmitted
Amazon	endogenous	endogenous	80%
eBay.8	endogenous	1	80%
eBay1	endogenous	1	100%
sealed bid	0	1	100%

Table 2. Experimental treatments

II. 2 Experimental procedure

The study was conducted with 30 groups of six participants each (8 groups each in Amazon, eBay.8, and eBay1, and 6 in sealed bid). Within each group we randomly rematched pairs of two bidders for a total of 20 auctions per bidder. A matching error in trial 19 rendered the data in all eBay.8 auctions numbers 19 and 20 incomparable.¹¹ Consequently, we report here only the results for trials 1-18 of *all* conditions in order to simplify the comparisons. We note, however, that there was no sign of an end game effect in any session, and that the conclusions we draw are invariant to whether we include auctions 19 and 20 in the sessions in which no problems occurred.

In each treatment, all rules were publicly explained with the help of example auctions (the instructions can be found in Appendix C). Auctions were run on networked computers,¹² and each participant could see on his screen both private and public information, updated after each period. The private information included the bidder's own private value, his own highest submitted bid so far, and a list of the auctions won earlier along with corresponding profits. The public information, known to both bidders, included the auction number (between 1 and 20), the period number within the current auction, the period type (stage 1 or stage 2), and the current price (at most an increment above second highest submitted bid). Participants were paid their cumulative earnings in the auctions plus a show-up fee of \$5 plus an additional \$5 if they were at least 5 minutes early.

III. Experimental Results

III.1 The evolution of late and early bidding

Figure 1 shows that the experimental results reproduce the main internet observations:¹³

- There is more late bidding in the fixed-deadline (eBay) conditions than in the automatic extension (Amazon) condition.
- Furthermore, as bidders gain experience, they are more likely to bid late in the eBay conditions, and less likely to bid late in the Amazon condition.

Figure 1a graphs the amount of late bidding (called “sniping” in the internet auction community) by recording the percentage of bidders who place a bid in stage 2.¹⁴ Since at most one stage 2-bid is recorded per bidder in each auction, these numbers can also be interpreted as the probability that a bidder will make a stage 2-bid. Each of the three multi-period auction conditions starts with about 40% of bidders submitting stage 2-bids, but by trial 18 Amazon has only about 10 percent, eBay.8 has 50 percent, and eBay1 has


¹¹ Some of those auctions had 3 bidders, while others had only 1.

¹² We used the z-Tree software toolbox by Fischbacher (1998).

¹³ Here and in the following, we will focus our analysis on the aggregate level (similar to the analysis of the field data in Ockenfels and Roth, 2001). Individual behavior is rather noisy, both for experienced and inexperienced bidders, which makes it difficult to identify and classify individual bidding patterns.

¹⁴ In Amazon, there may be more than one stage 2. In Figure 1 we included only the first stage 2 that determines whether the auction is extended at least once. We included also stage 2-bids that are lost, which


80 percent late bidders. We can reject the null-hypothesis that the overall numbers of stage 2-bids within each of the three auction types are from the same population (a Kruskal Wallis H -test based on the 8 independent sessions for each auction type yields $p = 0.000$). Overall, there are weakly significantly more stage 2-bids on eBay1 than on eBay.8 (two-sided Mann Whitney U -test, $p = 0.058$),¹⁵ and there are significantly more stage 2-bids in each of the eBay auction types, than on Amazon ($p = 0.000$, for each comparison separately). A probit analysis in Appendix A1 confirms the time trends are highly significant: in both eBay conditions the trend is towards more late bidding as bidders gain experience, while on Amazon experienced bidders submit fewer late bids.



(a) Percentage of bidders who place stage 2-bids (on Amazon: first stage 2), over time (by auction #)

happened with probability 0.2 on eBay.8 and Amazon. Figure 1a includes lines that show the results of simple OLS-regressions.

¹⁵ Restricting the analysis to experienced bidders (trials 10-18), the difference becomes significant at the 1 percent level.



(b) Number of stage 1-bids per bidder, over time (by auction #)

Figure 1. Number of bids per bidder and auction over time

Figure 1b graphs the number of stage 1-bids per bidder over time. Comparison of Figures 1a and Figure 1b shows that the rise in stage 2-bidding in the two eBay conditions is not part of a general increase in bidding activity, but just the opposite: the number of stage 1-bids is strongly decreasing in all three multi-period auctions.¹⁶

Overall, the average number of submitted bids in stage 1 and stage 2 (including lost stage 2-bids) per bidder and auction on Amazon, eBay.8 and eBay1 is 3.2, decreasing from 4.5 in trial 1 to 2.5 in trial 18. There are no statistically detectable differences in the bid numbers between the auction conditions (Kruskal Wallis H -test, $p=0.125$).

¹⁶ The OLS-regression in Appendix A2 confirms that the number of stage 1-bids is significantly decreasing with experience – particularly for the two eBay conditions. The statistical analysis also reveals that the number of a subject's stage 1-bids is increasing in the number of the opponent's stage 1-bids, implying that we see bidding wars in stage 1. These observations correspond to Ockenfels and Roth's (2001) field findings that the number of bids submitted by a bidder to an eBay auction is decreasing in experience as measured by his feedback rating, and increasing in the number of bids submitted by other bidders.

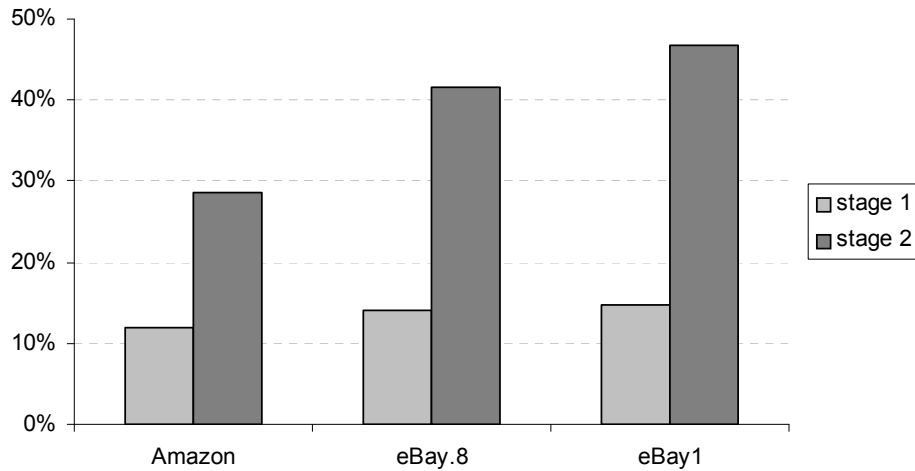


Figure 2. Share of bids submitted by current high bidder

Figure 2 shows that stage 1-bids are rarely placed by the current high bidder; early bids are mostly made in incremental bidding wars, when the low bidder raises his bid in an apparent attempt to gain the high bidder status. On the other hand, stage 2-bids in the eBay conditions are made almost equally often by the current high bidder and the current low bidder. That is, late bids on eBay appear to be planned by bidders regardless of their status at the end of stage 1.¹⁷

Our observations with respect to the bid timing not only replicate the field observations but also confirm game theoretic analyses. Ockenfels and Roth (2001) argue that there are multiple reasons why sniping may be a rational strategy (even) in a private value environment, in which there is a risk involved in late bidding. One intuition behind last-minute bidding at equilibrium is that there is an incentive to avoid a bidding war that raises the expected final price when there is still time for other bidders to react. Mutual delay until stage 2 can keep the final price down and therefore raise the expected profit of both bidders, because of the positive probability that another bidder's stage 2-bid will not be successfully transmitted in eBay.8. In Appendix B1 we elaborate on this idea in a

¹⁷ While there is no difference across auction types with respect to stage 1 (Kruskal Wallis H test, $p = 0.977$), there are significant differences with respect to stage 2 ($p = 0.000$). In particular, there are more high bidders submitting stage 2-bids in both eBay types, respectively, than on Amazon (two-sided Mann-

simplified model of implicit collusion.¹⁸ On Amazon, on the other hand, there is no way to delay one's bid until the opponent cannot react, because there is always time to respond to a successfully submitted bid. That is, the Amazon ending rule removes the advantage but not the risk of sniping, so that perfect Bayesian equilibrium bidding on Amazon does not involve stage 2-bids (see Appendix B2).

Sniping may also be a best response to incremental bidding that is observed both in the field (see Ockenfels and Roth, 2001 and 2002) and in our experimental setting. An incremental bidder starts with a bid below his value and is then prepared to raise his bid when he is outbid.¹⁹ Bidding late on eBay may be a best reply to incremental bidding, because this strategy would not give the incremental bidder any opportunity to respond to being outbid. In particular, by bidding in stage 2 in our eBay treatments, a bidder might win the auction against an incremental bidder, even when the incremental bidder's private value is higher. On Amazon, on the other hand, incremental bidders always have time to respond to the bidding activities of others, so that the incentive to snipe is eliminated (see Appendix B3 for a simple but more formal argument).

Comparing the timing of bids on eBay.8 and eBay1 allows us to assess the contributions to late bidding from implicit collusion by all bidders to avoid price wars, or from a best response by sophisticated bidders to incremental bidding by others. If

Whitney U test, $p = 0.001$ for each comparison separately), and there are more snipes by high bidders on eBay1 than on eBay.8 ($p = 0.045$).


¹⁸ We simplify the auction environment in the Appendix because the auctions in the experiment and on the internet are not exactly second-price auctions: the price is not exactly equal to the second highest bid, but is (at most) one discrete increment above it. The price increment creates incentives for the bidders to try to save (up to) one increment (25 cents in our experiments) by bidding just above the opponent's highest bid, which complicates the equilibrium analysis. (Recall that if the difference between the winning and the losing bid is smaller than the minimum increment, the price paid is equal to the highest bid. So the winner can save up to the increment by bidding slightly above the losing bid.) A much simpler theoretical treatment of our experimental environment is feasible if one abstracts away from the fact that the price exceeds the second highest submitted proxy bid by (at most) one minimum increment. This is how we proceed in the Appendix. In their models, Ockenfels and Roth (2001) take the price increment into account but restrict themselves to a particular distribution of values.

¹⁹ There are multiple reasons why bidders may want to bid incrementally in the field. For example, bidders can sometimes get information from others' bids that causes them to revise their valuations. Alternatively, increased attachment (such as the endowment effect) can yield higher valuations over time. None of these explanations is valid in our experimental environment, because values are exogenously induced and independent. Incremental bidding might also be caused by naïve, inexperienced bidders, who may be present both in the field and in our lab, and who mistakenly treat the eBay auctions like English first price auctions in which the high bidder pays his maximum bid. In fact, the field evidence in Ockenfels and Roth (2001) as well as the lab evidence in Figure 1b suggests that multiple bidding is negatively correlated with experience as measured by the feedback rating.


sniping is primarily occurring because of implicit collusion to keep prices down, we expected to see more stage 2-bids on eBay.8 than on eBay1, because the effect of late bidding on prices comes from the positive probability that the bid is lost. If, on the other hand, sniping is a reaction against incremental bidding, we expected to see the opposite, because a positive probability of a bid loss in this case reduces the expected benefit from sniping (see Appendix B3). The evidence shown in Figure 1a suggests that much of the sniping we see is a best response to incremental bidding, or at least that the late bidding in the eBay1 condition is not driven by rational collusion on the part of all bidders, since the benefits of implicit collusion by bidding late require $p < 1$.

III.2 The size of late and early bid increments, and price discovery


Figures 3a-c summarize by how much bids exceed the current minimum bid required, which is either the current price plus the minimum increment of \$0.25, or, before the first bid, the reservation price of \$1 (for the Amazon auctions, we graph all stage 2-bids, not only from the first stage 2). The graphs show the average increase of bids, conditional on bids being placed, so we have to interpret them together with the information in Figure 1, which shows the numbers of bids over time.



(a) Amazon



(b) eBay.8



(c) eBay1

Figure 3. Average increase of bids (conditioned on bidding) over current minimum bid

Figure 3a shows that the average size of the bid increments placed in stage 1 on Amazon clearly grows over time, while the size of bid increments in stage 2 does not reveal a strong time trend. This is consistent with our earlier observation that the *numbers* of both stage 1- and stage 2-bids in Amazon auctions decline over time. As bidders place fewer bids in stage 1, and hardly any bids in stage 2 (reflected by the large variances in Figure 3a), they bid in bigger increments in stage 1.

The situation is almost the opposite in each of the two eBay conditions. Figures 3b and 3c show that, while the average stage 1-bid increment stays relatively constant over time – slightly increasing on eBay.8 and slightly decreasing on eBay1 – the average stage 2-bid increment strongly grows in each eBay condition.²⁰ Moreover, these records tend to understate the difference between stage 1- and stage 2-bids over time because as Figure 1 showed, in both eBay conditions the stage 2-bid increments are getting larger at the same time as stage 2-bids are becoming more frequent, and stage 1-bids becoming less frequent.

As a result of these dynamics, the average stage 2-increment is about twice the size of stage 1-increments on eBay.8, and four times the size on eBay1, while it is only about half the size of stage 1-increments on Amazon. Mann-Whitney U tests based on the 8 independent sessions per auction type confirm these observations: There are no statistical differences between average bid increases across auction types with respect to stage 1, but the stage 2-increase is significantly larger in each of the eBay conditions than on Amazon ($p = 0.001$ for each comparison separately), and significantly larger on eBay1 than on eBay.8 ($p = 0.021$). Also, while on Amazon stage 1-increases are significantly larger than stage 2-increases ($p = 0.012$), the opposite is true in each of the eBay conditions ($p = 0.059$ and 0.002 for eBay.8 and eBay1, respectively). That is, as late bids become less frequent on Amazon they also become smaller, and as they become more frequent on eBay they also become larger. Thus, on eBay most of the ‘serious’ bidding is done in stage 2, while on Amazon most of the serious bidding is done in stage 1.

²⁰ Straightforward OLS regressions confirm all statements with respect to the time trends seen in Figure 3 at high significance levels.




Figure 4. Final stage 1-price (on Amazon: first stage 1) as percentage of final price and linear trends

This pattern of early and late bidding affects price discovery, that is, how well the price in the early part of the auction (at the end of stage 1 in both eBay conditions and at the end of the *first* stage 1 on Amazon) predicts the final price. Figure 4 shows that stage 1 prices are an increasingly good predictor for final prices on Amazon (after bidders gained experience, the stage 1-price reached more than 90 percent of the final price), whereas the opposite is true on eBay.8 (about 70 percent) and eBay1 (less than 50 percent).²¹

III.3 Learning how much to bid

Theoretically, all bids that exceed the private value are weakly dominated, regardless of the auction condition. So we expect most bids to be no higher than the induced private values.²² Furthermore, in all treatments with a definitely final period in stage 2 (sealed bid, eBay.8 and eBay1), we hypothesize that the final bids of most experienced bidders

²¹ The Spearman rank correlation coefficient between final (first) stage 1 price and final price is highest on Amazon (0.947), lower on eBay.8 (0.570), and lowest on eBay1 (0.270). All correlations are significant at $p = 0.000$.

²² Although we note that Kagel, Harstad, and Levin (1987) observed persistent overbidding by some bidders in second price auctions.

will be ‘close’ to their private values. To see why, first observe that bidding one’s value is not a dominant strategy, even in the sealed bid condition. This is because the minimum bid increment may create an incentive to bid an amount slightly above the opponents’ highest bid (but below the opponents’ highest bid plus the increment), so that the winner could avoid paying the entire minimum increment denoted by s .²³ However, since a winner can never advantageously affect the price in case of winning by bidding more than his value, and since a winner can never push the price down by more than one increment by bidding less, bidding one’s true value in the sealed bid auction may be called an “ s -dominant strategy.” That is, bidding one’s true value is a strategy that always yields a payoff not more than the minimum increment s below the maximum achievable by any other strategy, regardless of the strategies chosen by the other bidders. Since the truncated eBay-game that starts at stage 2 is a second-price sealed bid auction, an analogous argument holds for the eBay conditions. That is, any strategy that does not call for bidding value in stage 2 is ‘weakly s -dominated.’ Here, however, the minimum price increment may sometimes prevent bidders from bidding their exact values, and thus the final bids of experienced bidders can be expected to be within one increment below their private values in the eBay conditions.

On Amazon, on the other hand, the s -dominance criterion does not exclude final bids on Amazon that are substantially below value.²⁴ More importantly, a bidder on Amazon who is currently the high bidder has no incentive to increase his bid unless he is outbid, at which point he will always have the opportunity to raise his bid. So once he has exceeded the other bidder’s value, he has no incentive to increase his bid to his own value.

²³ Recall that the rules determine that the price can never exceed the highest submitted bid.

²⁴ To see this, suppose for instance that the opponent’s strategy in our Amazon condition is to bid \$1 in period 1 and then not to submit any more bids as long as the price does not exceed \$1.2, but to submit \$100 immediately after the price exceeded \$1.2. Then, facing this opponent, a bidder would earn zero by bidding his value, regardless of the timing, but could make a positive payoff by bidding, say, \$1.1 in period 1 yielding a final price of \$1.1.




Figure 5. Median of final bids (including lost stage 2-bids) as a percentage of value²⁵

Figure 5 shows that in all treatments the median of final bids relative to values is increasing over time and, as predicted, never exceeds 100 percent. But the bidding dynamics clearly differ across conditions. For inexperienced bidders, final bids in the sealed bid condition are substantially lower than final bids in the other conditions (up to trial 7). It appears that learning in the sealed bid auctions takes place across auctions, while learning in the dynamic auctions also takes place within auctions. For example, a bidder who imagines that he can win with a low bid does not learn that he is mistaken in a sealed bid auction until after the auction is over, but in the auctions conducted over time, he can revise his bid as soon as he is outbid.

For experienced bidders, Figure 5 shows that the medians of final bids in the eBay and sealed bid conditions converge to 100 percent of values.²⁶ On Amazon, on the other hand, the median bid of experienced bidders stays below 100 percent. This is consistent

²⁵ We show medians because there are a few outliers in one of the eight Amazon sessions between round 6 and 10 yielding high average bids in these rounds.

²⁶ The average proportion of bidders whose final bid is equal to the private value is lowest in the Amazon condition (0.125), somewhat higher in sealed bid (0.151), and much higher on eBay1 (0.234) and eBay.8 (0.251). In trial 18 the corresponding numbers are 0.167, 0.111, 0.313, and 0.333. So, the fact that final bids approach values in aggregate (Figure 5) does not imply that all experienced bidders actually bid

with the theoretical considerations explained above. Furthermore, note that incremental bidders learn on eBay that they are sometimes outbid in stage 2 at prices more than an increment below values, which conceivably lead them to bid closer to values over time. Incremental bidders on Amazon, on the other hand, are never outbid at prices more than an increment below their values, regardless of how their final bids relate to the values.²⁷ Thus, for incremental bidders, the pressure to learn to bid one's value is weaker on Amazon than on eBay. Once incremental bidding has reached the second highest value, the high value bidder has no incentive to bid up to his own value.

III. 4 Revenue and efficiency

Figure 6 shows the efficiency across all conditions measured as the average frequency the auctions are won by the bidder with the higher value, and Figure 7 shows median revenues.




Figure 6. Average efficiency

values. Conversely, of course, the observation that many bidders are not bidding their values does not imply that behavior is far from equilibrium.

²⁷ This is the reason why sniping is not a best response to incremental bidding on Amazon.

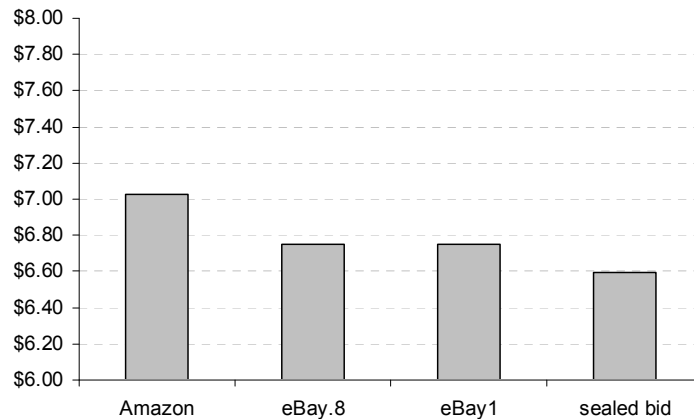


Figure 7. Median revenues

The Amazon condition is slightly more efficient and yields higher revenues than the other conditions (applying a one-sided Mann-Whitney U test, all pairwise comparisons with Amazon yield significance at the 6 percent level, for efficiency and revenue separately). On the other hand, revenues and efficiency are lowest in the sealed bid treatment (all comparisons are significant at the 6 percent level). This seems to reflect the fact that Amazon is the only treatment where low bidders always had time to respond to being outbid at prices below values, eBay-bidders could only respond to stage 1-bids but not to stage 2-bids, and losers in sealed bid never had the opportunity to respond to the bids of other bidders.²⁸

The risk involved in stage 2-bids in the eBay.8 condition significantly reduces efficiency and revenues for experienced bidders (who snipe most often). Restricting the analysis to rounds 10 to 18, average efficiency is 98 percent on eBay1 and 88 percent on

²⁸ The efficiency/revenue pattern across Amazon and eBay.8 is consistent with our observation from Section III.1 that the interaction between ‘naïve’ incremental and ‘sophisticated’ bidders (who play best response against incremental bidding) affects the results. Incremental bidders may win Amazon auctions with bid amounts that are substantially below their values. At the same time, as mentioned in Section III.3, on Amazon there is no pressure on incremental bidders to learn to bid values. Thus, average bids of experienced bidders on Amazon can be expected to be lower than eBay.8-bids (as supported by Figure 5). However, this does not necessarily imply that prices and revenues are also lower: Prices on Amazon can be expected to be close to the second highest value, because both incremental and sophisticated bidders are prepared to bid up to their values, and they always have time to do so (but unlike sophisticated bidders, incremental bidders will not bid more than they have to in order to win). Thus, if there is both sophisticated and naïve bidding on Amazon, the bidder with the higher value wins at a price close to the second highest bid, even though average bids will be below values. Incremental bidders on eBay.8, on the other hand, learn to bid values over time (see Figures 1b and 5). But since more bids are lost on eBay.8 than on Amazon (see Figure 1a), both efficiency and revenue can be expected to be smaller.

eBay.8, and average revenues are \$6.73 on eBay1 and \$6.62 on eBay.8. When aggregating over all rounds, however, efficiency and revenues in the eBay conditions are statistically indistinguishable

IV. Conclusions

The experiment presented here was designed to investigate the effects of auction closing rules on bidding behavior. It was motivated by comparisons of bidder behavior on eBay and Amazon. The experiment confirms, under controlled conditions, that the difference in ending rules between eBay and Amazon is sufficient to cause the patterns of behavior observed in the field data. The results show that there is much more late bidding with the eBay fixed ending rule than with the Amazon automatic extension rule, and that this tendency increases with experience. The data also show considerable incremental bidding, which contributes to the different late bidding patterns on eBay and Amazon. With experience, incremental bidding is reduced but not eliminated, as was observed in the field data.

The experiment also allowed us to observe aspects of behavior that are not readily available in the field data. In particular, our data suggest that, on average, bidders learn to bid their values, though in keeping with the strategic incentives, Amazon-bids are slightly below values even for experienced winning bidders. Still, in our experimental environments, Amazon is the format with the highest efficiency and revenues, because the incentive for strategic delay of bids is low. One possible reason for the efficiency of the Amazon condition could be due to the learning opportunities within the Amazon conditions, which are better than the learning opportunities within the eBay and the sealed bid conditions.

Note that, despite the superior control that we achieve in the laboratory, if we presented *only* experimental data we could not be confident that the same effect would be observed on the internet. It might be that, in the laboratory, people bid late because it gives a slight advantage and has little cost: as they are already committed to staying until the end of the experiment. In real life, it might be supposed that people have better things to do. The fact that we see the same pattern of behavior both in the lab and in the field gives us an indication of its robustness.

While the results of the experiment replicate the basic observations in the field data, we do not claim that the field data are *fully explained* by the experimental data. By design, the experimental setting eliminated many of the factors present in the field data (number of bidders, multiple items offered simultaneously, heterogeneity of sellers, bidders, and products, etc). By eliminating these sources of variation, the experiment showed that they are not necessary to produce the observed differences: the difference in auction rules is sufficient. However, that is not to say that none of the factors that we eliminated from the experiment could not nevertheless contribute to the effects observed in the field data. For example, while the experimental results show that we get the predicted effect even when we control for number of bidders, that is not to say that the number of bidders does not have an effect on how bidding compares on eBay and Amazon. The experimental results also do not tell us whether these different auction formats would attract different numbers of buyers and sellers if they were free to self select, as in the field data. That is, the higher Amazon revenues we observe in the experiment, holding the number of bidders constant, might attract sellers to choose automatic extensions, but maybe the prospect of higher bidder profits on eBay.⁸ would attract additional bidders, which would change sellers' choices, etc. The experimental data demonstrate some sufficient conditions for late bidding, but not necessarily the full set of factors that take place on the internet.

Thus the experimental and field data, together with the theory developed to explain them, are complements, not substitutes. Together they help us to understand how, in auctions as well as in other markets, the rules of the market influence the timing of transactions, which can have important implications for prices and efficiency.

References

- Asker, John, Brit Grosskopf, Carl N. McKinney, Muriel Niederle, Alvin E. Roth, and Georg Weizsacker. "Teaching Auction Strategy Using Experiments Administered via the Internet." Working paper, Harvard University, 2002.
- Avery, Christopher, Christine Jolls, Richard A. Posner, and Alvin E. Roth. "The Market for Federal Judicial Law Clerks." *University of Chicago Law Review*, 68, 3, Summer 2001, 793-902.
- Bajari, Patrick and Ali Hortaçsu. "Winner's Curse, Reserve Prices and Endogenous Entry: Empirical Insights from eBay Auctions." Working paper, Stanford University, 2000.
- Bolton, Gary, Elena Katok, and Axel Ockenfels. "How Effective are Online Reputation Mechanisms? An Experimental Investigation." Working paper, Harvard Business School, 2002.
- Ceci, Stephen J., and Urie Bronfenbrenner. "Don't forget to take the cupcakes out of the oven: Prospective memory, strategic time-monitoring, and context." *Child-Development*, 56(1), February 1985, 152-164
- Croson, Rachel. "Partners and Revisited." *Economics Letters*, 1996, 53, 25-32.
- Dellarocas, Chrysanthos. "Efficiency and Robustness of Mediated Online Feedback Mechanisms: The Case of eBay." Working paper, Sloan School of Management, MIT, June 2002.
- Dews, Peter B. "The effects on the pattern of responding of changes in requirements at reinforcement." *Journal of the Experimental Analysis of Behavior*, 1969, 12, 191-199.
- Feltovich, Nick. "Reinforcement-Based vs. Beliefs-Based Learning in Experimental Asymmetric-Information Games," *Econometrica*, 68 (2000), 605-641.
- Ferster, C. B and Skinner B F (1957) *Schedules of Reinforcement* Appleton-Century-Crofts, Inc. New York.
- Fischbacher, Urs. "z-Tree. Zurich Toolbox for Readymade Economic Experiments." Working Paper, University of Zurich, 1998.
- Gächter, Simon, and Arno Riedl. "Moral Property Rights in Bargaining." Working paper, CREED University of Amsterdam, 2002.
- Güth, Werner, Maria Vittoria Levati, and Boris Maciejovsky. "Deadline Effects in Ultimatum Bargaining: An Experimental Study of Concession Sniping with Low or

- no Costs of Delay.” Working Paper, Max Planck Institute for Research into Economic Systems, Jena, 2001.
- Haruvy, Ernan, Alvin E. Roth, and M. Utku Ünver. “The Dynamics of Law Clerk Matching: An Experimental and Computational Investigation of Proposals for Reform of the Market.” Working paper, Harvard University, October 2001.
- Kagel, John H. , Ronald H. Harstad, and Dan Levin. “Information Impact and Allocation Rules in Auctions with Affiliated Private Values: A Laboratory Study.” *Econometrica*, 55, 1987, 1275-1304.
- Kagel, John H. and A.E. Roth. "The Dynamics of Reorganization in Matching Markets: A Laboratory Experiment Motivated by a Natural Experiment." *Quarterly Journal of Economics*, February 2000, 201-235.
- Malhotra, Deepak, and J. Keith Murnighan. “Milked for all Their Worth: Competitive Arousal and Escalation in the Chicago Cow Auctions.” “Working paper, Kellogg School of Management, Northwestern University, April 2000.
- Milgrom, Paul. “*Auction Theory for Privatization*.” In preparation, 2002.
- Ockenfels, Axel and Alvin E. Roth. “Late and Multiple Bidding in Second Price Internet Auctions: Theory and Evidence Concerning Different Rules for Ending an Auction.” Working paper, Harvard University, 2001.
- Ockenfels, Axel and Alvin E. Roth. “The Timing of Bids in Internet Auctions: Market Design, Bidder Behavior, and Artificial Agents.” *Artificial Intelligence Magazine*, Fall 2002, 79-87.
- Resnick, Paul and Zeckhauser, Richard. “Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay’s Reputation System.” In: Michael R. Baye (editor), *The Economics of the Internet and E-Commerce*, Advances in Applied Microeconomics, 11, 2002, Amsterdam, Elsevier Science.
- Roth, Alvin E. “The Economist as Engineer: Game Theory, Experimental Economics and Computation as Tools of Design Economics.” Fischer Schultz lecture, *Econometrica*, 70, 4, July 2002, 1341-1378.
- Roth, Alvin E., and Ido Erev. "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term." *Games and Economic Behavior*, Special Issue: Nobel Symposium, 8, January 1995, 164-212.
- Roth, Alvin E., Murnighan, J.K., and Schoumaker, F. "The Deadline Effect in Bargaining: Some Experimental Evidence." *American Economic Review*, September 1988, 78(4), 806-23.

- Roth, Alvin E., and Axel Ockenfels. "Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet." *American Economic Review*, September 2002, 92(4), 1093-1103.
- Roth, Alvin E. and Peranson, Elliott. "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design." *American Economic Review*, 89(4), September 1999, 748-780.
- Roth, A.E. and X. Xing. "Jumping the Gun: Imperfections and Institutions Related to the Timing of Market Transactions." *American Economic Review*, 84, September 1994, 992-1044.
- Wilcox, Ronald T. "Experts and Amateurs: The Role of Experience in Internet Auctions." *Marketing Letters*, 2000, 11(4), 363-374.
- Vickrey, William (1961). "Counterspeculation, Auctions, and Competitive Sealed Tenders." *Journal of Finance*, 16, 8-37.

Appendix A: Statistics

A1. Number of stage 2-bids

<i>Random Effects Probit Model</i>	
Maximum Likelihood estimates (and <i>T</i> -statistics)	
Dependent variable = “1” for stage 2-bid (on Amazon: <i>first</i> stage 2) per bidder and per auction, and “0” else	
<i>Independent variables</i>	<i>Coefficients</i>
Constant ^a	- 0.287** (- 2.927)
Trial number (between 1 and 18) if Amazon and 0 else	- 0.101** (- 12.249)
Trial number if eBay.8 and 0 else	0.020** (3.575)
Trial number if eBay1 and 0 else	0.104** (13.996)
Rho ^c (random effects)	0.486** (13.161)
Number of observations ^d	2592
Log-likelihood	- 1184.896

* denotes significance at the 5 percent level (two-sided), ** significance at the 1 percent level (two-sided).

^a There is no statistically significant level effect across treatments, so we do not include treatment dummies here.

^b All three pairwise comparisons of the effects of trial numbers across treatments yield significance at the 1 percent level.

^c Individual subject differences are clearly present as indicated by the highly significant RHO, the Hausman test statistic for the presence of random effects.

^d Each bidder in each auction is one observation, making a total of 2592 observations (= 3 treatments*48 bidders per treatment*18 trials).

Table A1. Probit regression: Stage 2-bids

The regression shows that over time the frequencies of sniping decrease on Amazon, increase on eBay.8 and even stronger increase on eBay1.

A2. Number of stage 1-bids

OLS estimates (and <i>T</i> -statistics)	
Dependent variable = number of stage 1-bids per bidder and auction	
<i>Independent variables</i>	<i>Coefficients</i>
Constant ^a	2.752** (22.348)
Trial number (between 1 and 18) ^a	- 0.073** (- 8.480)
Number of stage 1-bids by the opponent ^a	0.335** (17.586)
Number of own stage 2-bids if Amazon, zero else	0.351** (2.827)
Number of own stage 2-bids if eBay.8 (0 or 1), zero else	- 0.354** (- 2.907)
Number of own stage 2-bids if eBay1 (0 or 1), zero else	- 0.874** (- 8.003)
Number of observations ^b	2592
<i>R</i> ²	0.205

* denotes significance at the 5 percent level (two-sided), ** significance at the 1 percent level (two-sided).

^a There is no statistically significant level effect across treatments, so we do not include treatment-specific variables here.

^b Each bidder in each auction is one observation, making a total of 2592 observations (= 3 treatments*48 bidders per treatment*18 trials).

Table A2. OLS regression: Stage 1-bids

The regression shows that the number of stage 1-bids decreases over time, that the opponent’s bidding activity seems to trigger ‘reciprocal’ stage 1-bidding, and that if bidders in the eBay conditions decide to bid in stage 2 they bid less in stage 1, while on Amazon stage 2-bids lead to even more stage 1-bids during the extension of the auction.

Appendix B: Theoretical considerations in a simplified model

As explained in footnote 18 of the main text, the following analyses simplify the eBay and Amazon environments by abstracting away from the fact that the price exceeds the second highest submitted proxy bid by (at most) one minimum increment. All other details of the auction formats, including the minimum bid increment requirement, are taken into account.²⁹

B1. Late bidding equilibrium on eBay.8

Proposition [eBay.8]: In the simplified eBay.8 experimental environment, there exists a perfect Bayesian equilibrium in undominated strategies in which no bids are placed until stage 2, at which time bidders bid their private values.

Proof: There exist multiple equilibria on eBay.8. In particular, by applying Vickrey's (1961) argument analogously to our (simplified) model, it is easy to see that there is an equilibrium in which all bidders bid their values in period 1 of stage 1 and then do not bid anymore until the auction is over.³⁰ But there are also equilibria in which both bidders submit values in stage 2 and do not bid in stage 1, even though stage 2-bidding involves a risk that the bid is lost.

Extending the example of Ockenfels and Roth (2001),³¹ consider the following late bidding strategies, which we will show constitute an equilibrium for risk neutral bidders on eBay.8. On the equilibrium path, each bidder i 's 'sniping strategy' is not to bid until stage 2 and then to bid his value, unless the other bidder deviates from this strategy by bidding in stage 1. Off the equilibrium path, if player j places a bid in period 1 of stage 1, then player i bids his true value in period 2 of stage 1. That is, each player's strategy is to do nothing until stage 2, unless the other bidder makes a stage 1-bid, that would start a

²⁹ That is, analogous to eBay and Amazon each new bid must be a minimum increment above the current price, but unlike the internet auction counterparts we assume here that the current price always exactly equals the second highest submitted bid.

³⁰ This is the kind of equilibrium behavior eBay promotes when it explains why it recommends "bidding the absolute maximum that one is willing to pay for an item early in the auction" on its auction sites (<http://pages.ebay.com/aw/notabuse.html>, 1999).

³¹ In their example, all bidders had the same private value and this was common knowledge among all players.

price war at which the equilibrium calls for a player to respond by bidding his true value in the subsequent period.

Suppose for the moment that bidder 1's value is \$10, the highest possible value in our experiment, and bidder 2's value is \$6, the smallest possible value in our experiment. Let $p = 0.8$ be the probability of a successfully transmitted bid at stage 2, as on eBay.⁸ If bidders follow the strategy described above, bidder 1 earns \$9 (= value – minimum bid) if his bid at stage 2 is successfully transmitted and the other bidder's bid is lost, which happens with probability 0.16 (= $p(1 - p)$), and he earns \$4 (= value – opponent's value) if both stage 2-bids are successfully transmitted, which happens with probability 0.64 (= pp). If bidder 1's bid is lost at stage 2, which happens with probability 0.2 (= $1 - p$), his payoff is zero yielding a total expected payoff of \$4 for bidder 1. Similarly, bidder 2 earns \$5 (= value – minimum bid) if his stage 2-bid is successfully transmitted and the opponent's stage 2-bid is lost, which happens with probability 0.16, and zero otherwise, yielding a total expected payoff of \$0.80 for bidder 2.

Unilateral deviation from the sniping strategy is not profitable for either bidder. First, in stage 2, any other bid than the true value is part of a weakly dominated strategy. (Recall that in our simplified environment stage 2 is Vickrey's second price sealed bid auction.) Second, in stage 1, any bid triggers an 'early' price war in which each player bids his true value in stage 1 (which constitutes a Nash equilibrium in our model). The price war yields a payoff of \$4 (\$10 – \$6) for bidder 1 and zero for bidder 2, which is equal to the corresponding sniping payoffs for bidder 1, and which is smaller than the corresponding sniping payoff for bidder 2. This proves that the sniping strategy is a best reply for bidders with values \$10 and \$6.

In fact, the sniping strategies constitute an equilibrium for any realizations of the values. To see why, observe that for a bidder 1 with value $v_1 > v_2$ the expected profit from mutual sniping is $0.16*(v_1 - \$1) + 0.64*(v_1 - v_2)$, while the expected payoff after an early bidding war (that is, after mutually bidding true values in stage 1) is $v_1 - v_2$. Inspection shows that the difference of these payoffs (= $-0.2v_1 + 0.36v_2 - 0.16$), is decreasing in v_1 and increasing in v_2 . In the last paragraph, we have shown that if v_1 takes the maximal value (\$10) and v_2 takes the minimal value (\$6), the sniping strategies constitute an equilibrium. Hence, all other combinations of private values make sniping

even more profitable for a bidder 1 with $v_1 > v_2$. Since the sniping strategies always yield a higher payoff for the bidder with the lower value compared to an early bidding war, the sniping strategies constitute an equilibrium for all combinations of private values.

B2. The Amazon case

On Amazon there is no way to delay one's bid until the opponent cannot react, because there is always time to respond to a successfully submitted bid. That is, the Amazon ending rule removes the advantage but not the risk of sniping. Consequently, under very mild additional assumptions (to deal with cases in which players are indifferent between bidding and not bidding) perfect Bayesian equilibrium bidding on Amazon cannot involve stage 2-bids. Specifically, we assume a 'willingness to bid' in that a bidder prefers to earn zero by bidding and winning the auction rather than by not submitting a bid (and hence earning zero).³²

Proposition [Amazon]: Assuming a willingness to bid, there are no stage 2-bids at a perfect Bayesian equilibrium in undominated strategies in our simplified Amazon model.

Proof. We extend the example of Ockenfels and Roth (2001).³³ At a perfect Bayesian equilibrium in undominated strategies:

1. No bidder ever bids above his value: Any strategy that calls for bidder j to bid above v_j in any period t is dominated by the otherwise identical strategy in which j bids at most v_j at period t .
2. There is a finite period t^* such that the auction receives its last bids by period t^* , because proxy bids must rise by at least 25 cents with each new submission and because no bidder will ever submit a reservation price greater than $v_{\max} = \$10$. If the

³² The 'willingness to bid' assumption would not alter our eBay.8 proposition. It is a weak assumption on preferences, since it only comes into play when bidders are indifferent. But since we need this weak assumption, this is a good point to warn against over-interpreting the theorem. Different reasonable assumptions (e.g. allowing imperfect equilibria) can yield somewhat different conclusions. The point of the theorem, however, is that the incentives for late bidding in Amazon are very different from those on eBay: this assumption leaves late bidding equilibria intact on eBay, but rules them out on Amazon.

auction gets to this period, there is only room for the price to rise by no more than 25 cents.

3. In principle, the last period t^* with bidding activity may either be a stage 1- or a stage 2-period. However, the bidder who at t^* is not the current high bidder and who has a value greater by 25 cents than the current price will – and by our experimental design can – make sure that t^* is a stage 1-period so that his last bid is transmitted with certainty (recall that a stage 2-period can only be reached if no bid is submitted in the previous period). Here, the willingness to bid comes into play, because it rules out possible indifference between bidding and not. Since no bidder is indifferent between casting the winning bid and not, any strategy profile that caused a player to bid at a stage 2-period would have a lower expected payoff (because $p < 1$) than a strategy at which he bid at stage 1, when bids are submitted with certainty. And, because this will be the last period with bids in the auction, the standard Vickrey second-price private value argument implies that a bid of less than the true value would constitute part of a dominated strategy: it could only cause some profitable opportunities to be missed.
4. Inductive step. Suppose at some period t , it is known that at the *next* period the bidders who are not the current high bidder and who have a value greater by at least 25 cents than the current price will place bids in the amount of their values with certainty. Then all bidders will bid their true values in a stage 1 period. Since a price war will result if the auction is extended by a successful bid at a stage 2-period, any strategy profile that calls for a bidder who is not already the high bidder to bid at stage 2 is not part of an equilibrium, since that bidder gets a higher expected return by bidding his true value at a stage 1-period. As a result, there are no stage 2-bids in any perfect Bayesian equilibrium.

B3. Incremental bidding

An incremental bidder does not use the ‘proxy bidding agent’ but starts with a bid below his value and is then prepared to raise his bid whenever he is outbid. Bidding late

³³ As in the eBay case, their example is characterized by identical private values.

may be a best reply to incremental bidding, because bidding very near the deadline of an auction with a hard close would not give the incremental bidder an opportunity to respond to being outbid. In the following we will for simplicity restrict ourselves to a straightforward (naïve) form of incremental bidding defined as a strategy that calls for bidding in minimum increments until the high bidder status is reached, but not more than the private value.

Proposition [Incremental bidding]: The gain from bidding in stage 2 (“sniping”) against an incremental bidder in our simplified environment is always strictly positive on eBay1 and strictly positive but smaller on eBay.8. Sniping is not a best response against incremental bidders on Amazon.

Proof: Let us start with eBay1 and suppose bidder j knows that he is matched with an incremental bidder i . If j refrains from bidding early and bids his value in stage 2, he will win the auction for sure at a price of \$1, because i bids \$1 in the first period (which is the smallest bid sufficient to reach the high bidder status) and then never again, since he only realizes that he was outbid in stage 2 when the auction is over. On the other hand, each bid by bidder j in stage 1 increases the final price to at least \$1.25. Consequently, bidding late is always a best response against an incremental bidder on eBay1.

The eBay.8 case is more complicated since late bids get lost with positive probability which creates a cost of sniping. Suppose bidder j with value v_j knows that he is facing an incremental bidder i . If j bids his value in stage 1, his profits are positive if and only if $v_j \geq v_i$, where v_i denotes the incremental bidder’s value. The expected payoff from this strategy is \$1. If, on the other hand, bidder j bids late, he wins with probability 0.8 at a price of \$1. Inspection shows that it is always (that is, for all values v_j) advantageous not to get involved in an early price war with an incremental bidder. However, the incentive to refrain from bidding early is smaller than on eBay1 since the risk of late bidding reduces the expected benefit from waiting until stage 2.

The Amazon case is trivial. Any late bid either extends the auction so that an incremental bidder can respond with probability one, or it is lost. An early bid also

extends the auction so that an incremental bidder can respond with probability one, but it is transmitted with certainty.

Admittedly, the ‘pure’ form of incremental bidding as assumed for the Theorem is rarely observed in our experiments. Rather, incremental bidding typically involved bidding in larger than minimum increments, and also did not exclude the possibility of a stage 2-bid. However, in principle the mechanics of the Theorem are valid for other kinds of incremental bidding patterns as long as ‘incremental’ bids are provoked in response to early bids and thus drive up both the early and the final price.

Appendix C: Instructions

This appendix contains the written instructions (in small font) that were given to the subjects for all sessions of the experiment. The monitor read all instructions aloud. Then, a public demonstration on large screen explained the computer interface of the auctions by way of example auctions. After the demonstration and throughout the sessions, subjects were given the opportunity to ask clarifying questions. Answers to individual subject questions were given in private. Subjects were told not to communicate verbally with other participants. During the experiment subjects sat in front of a partitioned computer (so that they could not see the activity of other participants) until the end of the experiment. Once the experiment was over, subjects were thanked, paid according to their performance, and dismissed.

General instructions as handed out to all subjects in all sessions

The purpose of this session is to study decision making in an auction market. In this study you will participate in a series of auctions. In each auction, there will be exactly two bidders, you and one other. You will not be paired with the same bidder from one auction to the next; you will always bid against one different bidder. All bidders receive exactly the same instructions.

What is auctioned and what are my payoffs?

The objects for which you will be bidding in these auctions will be artificial commodities. You will be informed at the outset of the auction how much the object for sale is worth to you. This sum of money, called your personal value, v , will be randomly selected from a uniform distribution between \$6 and \$10. At the end of the auction, if you won the auction and paid a price p , you will have earned $v - p$. If you do not win the auction, you will earn nothing for that auction. If the difference is negative, it represents a loss. Gains are added to and losses are subtracted from the show-up fee and paid to you in cash at the end of the experiment.

For example:

Your personal value	Amount paid	Your earnings
\$7	\$6	\$1
\$8	\$6	\$2
\$5	\$6	-\$1

How can I bid?

The auctions you will participate in today are ‘second price auctions.’ This means that the price paid is the second highest bid plus at most 25¢ (you will never have to pay more than your bid). In other words, if you are the highest bidder for one of the objects being auctioned, the amount you will pay will be set at the value of the second highest bid plus at most 25¢. If both bidders submit the same exact bid, earlier bids of the same amount take precedence. In case of simultaneously submitted bids of the same amount, one of the bids is chosen randomly to be the earlier bid.

For example:

Bidder 1's highest bid	Bidder 2's highest bid	Highest bidder	Price paid
\$10	\$6	1	\$6.25
\$9	\$6	1	\$6.25
\$6.10	\$6	1	\$6.10
\$6	\$6	1 or 2	\$6

So, your task is to submit bids for the item in competition with the other bidder. Because of the second price auction rules, the bid is not necessarily equal to the price but rather an upper bidding limit. Once the bid is submitted, the software (working as a proxy for you) will determine the price by bidding only enough to outbid the other bidders as the auction proceeds. In other words, you can submit a bid and have the resulting price (at most) 25¢ above the next-highest competing bid. This continues until someone exceeds your bid or the auction ends.

- While your bid is kept secret, the price is public information.
- Whenever you submit a bid, it must satisfy three conditions.³⁴
 - First, it must be greater than or equal to the minimum bid of \$1.
 - Second, it must exceed the current price by at least 25¢.
 - Third, it must exceed the bid you placed before (if any).

If one or more conditions do not hold, your bid will not be accepted and you will be asked to submit another one.

Session-specific instructions for sealed-bid

In each auction you will be given one and only one opportunity to submit your bid for each product. Once all the participants submit their bids, the highest bidder will win the auction for the price set by the second highest bidder, plus (at most) 25¢. The winner will earn $v - p$ for that auction. At that point the next auction will start.

Session-specific instructions for eBay1

In each auction there will be two stages. Every auction will start with the ‘regular bidding stage’. During this stage, all bidders will have the opportunity to submit their bids and once they all submit their bids, a new round will start giving you information from the previous round (who is the highest bidder and at what price). The regular bidding stage will continue until there is a round in which no bidder increases his or her bid. At that point, the auction will move to the ‘final bidding stage’. This stage is very similar to the regular bidding stage but with one important difference: There is only one round and once this round is over, the auction is finished.

³⁴ In the sealed bid condition it was indicated that there is only one requirement, which is bid at least \$1. The other requirements do not apply and were thus not mentioned in the instructions for this condition.

Once all the participants submit their bids in the final bidding stage, the highest bidder will win the auction for the price set by the second highest bidder plus (at most) 25¢. The winner will earn $\$v - p$ for that auction. At that point the next auction will start.

Session-specific instructions for eBay.8

In each auction there will be two stages. Every auction will start with the ‘regular bidding stage’. During this stage, all bidders will have the opportunity to submit their bids and once they all submit their bids, a new round will start giving you information from the previous round (who is the highest bidder and at what price). The regular bidding stage will continue until there is a round in which no bidder increases his or her bid. At that point, the auction will move to the ‘final bidding stage’. This stage is very similar to the regular bidding stage but with two important differences:

1. There is only one round and once this round is over, the auction is finished.
2. The probability that any bid will be successfully transmitted is reduced from 100% in the regular bidding stage to 80% in the final bidding stage. (That is, there is a 20% chance that a bid submitted in the final bidding stage will not be received by the auctioneer.)

Once all the participants submit their bids in the final bidding stage, the highest bidder whose bid is received by the auctioneer in either stage will win the auction for the price set by the second highest bidder whose bid is received in either stage, plus (at most) 25¢. Bids that are not received by the auctioneer will not influence the final outcome. The winner will earn $\$v - p$ for that auction. At that point the next auction will start.

Session-specific instructions for Amazon

In each auction there will be at least two stages. Every auction will start with the ‘regular bidding stage’. During this stage, all bidders will have the opportunity to submit their bids and once they all submit their bids, a new round will start giving you information from the previous round (who is the highest bidder and at what price). The regular bidding stage will continue until there is a round in which no bidder increases his or her bid. At that point, the auction will move to the ‘second bidding stage’. This stage is very similar to the regular bidding stage but with two important differences:

1. There is only one round and once this round is over, the auction is either finished, or resumed in the regular bidding stage.
2. The probability that any bid will be successfully transmitted is reduced from 100% in the regular bidding stage to 80% in the second bidding stage. (That is, there is a 20% chance that a bid submitted in the final bidding stage will not be received by the auctioneer.)

If there are no successful bids during the second bidding stage (either because no bidder increased his or her bid or because no bids were received by the auctioneer), the auction will end. However, if there is at least one successful bid during the second bidding stage, the auction will resume in the first bidding stage and will advance again to the second bidding stage only after a round in which no bids are submitted (and so forth). Note that as long as there are successful bids during the second bidding stage, the auction will not end.

Once all the participants submit their bids in the second bidding stage, the highest bidder whose bid is received by the auctioneer in either stage will win the auction for the price set by the second highest bidder whose bid is received in either stage, plus (at most) 25¢. Bids that are not received by the auctioneer will not influence the final outcome. The winner will earn $\$v - p$ for that auction. At that point the next auction will start.