An experimental analysis of ending rules in Internet auctions

Dan Ariely[*]

Axel Ockenfels[**]

and

Alvin E. Roth[***]

A great deal of late bidding has been observed on eBay, which employs a second price auction with a fixed deadline. Much less late bidding has been observed on Amazon, which can only end when ten minutes have passed without a bid. In controlled experiments, we find that the difference in the ending rules is sufficient by itself to produce the differences in late bidding observed in the field data. The data also allow us to examine bid amounts in relation to private values, and how behavior is shaped by the different opportunities for learning provided in the auction conditions.

1. Introduction

■   How to end an auction is a subject of active interest 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 the auctions run on eBay and Amazon besides their rules, which potentially makes it difficult to attribute the differences in observed behavior to the different ending rules. Here we present the results of a laboratory experiment designed to complement the field data by controlling away all differences except those in the auction rules, in order to unambiguously test how the rules contribute to differences in bidding behavior.

The main 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, on the other hand, are automatically extended if necessary, past the scheduled end time, until ten minutes have passed without any bid having been submitted.

    [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, 2004; Roth, 2002).

892 / THE RAND JOURNAL OF ECONOMICS

Roth and Ockenfels (2002) compare the timing of bids in eBay and Amazon second-price auctions on the Internet, and observe that bids in eBay auctions are much more concentrated near the end of the auction than bids in Amazon auctions. Furthermore, 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% of the Amazon auctions received bids (i.e., only 11% of the Amazon auctions were extended past the scheduled deadline), while more than half of the eBay auctions received bids in the last ten minutes (and over 10% of the eBay auctions received bids in the last ten seconds).

Surveys of late bidders by Roth and Ockenfels (2002) showed that there are two sources of risk involved in late bidding (also known as "sniping") on eBay. 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. The survey responses suggested that, on average, there is a 20% risk that one of these two reasons lead to a planned late bid not being successfully transmitted.[2] This risk is reduced but not eliminated by the use of artificial bidding agents. eSnipe.com 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 reported, on the basis of more than 4,200 bids per day, that on average 4.5% of eSnipe's bids failed to be successfully transmitted in September 2000 (www.esnipe.com/stats.asp.2000). One of the variables in our experiment will be the probability that a late bid fails to be transmitted to the auction.

Ockenfels and Roth (forthcoming) demonstrate that there are multiple reasons that can contribute to why the difference in ending rules between eBay and Amazon can 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. Ockenfels and Roth (forthcoming) show that in such an environment, late bidding can arise as a rational response to many different causes. In both private-value and common-value auctions, and both in equilibrium and as a 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 multiple 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 which auctions to participate in, 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 (see Roth and Ockenfels, 2002). Some combination of these uncontrolled differences between eBay and Amazon might in fact be the cause of the observed difference in bidding behavior, instead of the differences in rules.

Also the proxies for experience in the field data ("feedback ratings") used by Ockenfels and Roth (forthcoming) are imperfect. For example, feedback ratings reflect only 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 have not only 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, they may have lower opportunity cost of time and thus can spend the time to bid late, or they may be more willing to pay the fixed cost of purchasing and learning to use a sniping program.[3]

---

[2] In Roth and Ockenfels' survey, fewer than 10% of the late bidding resondents reported that they had used automated bidding software.

[3] For further discussion of eBay's rating, see Dellarocas (forthcoming), Resnick and Zeckhauser (2002), and Bolton, Katok, and Ockenfels (2004).

Although 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 observed differences in bidding behavior on eBay and Amazon. For instance, the fact that bids on eBay antique auctions are even more skewed toward 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 (e.g., Bajari and Hortaçsu, 2003; Ockenfels and Roth, forthcoming). 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 simplest case of purely private values.[4]

Here 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 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 that help us not only to investigate how the auction ending rule contributes to late bidding, but also to identify different factors that contribute to late bidding.[5] The experimental data will also allow us to compare the revenues and relative efficiency resulting from the different types of auctions.

2. Experimental environment

■   **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 (80% in eBay.8 and 100% in eBay1). 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 $10. The winner of an auction received his private value minus the final price, and a loser received nothing for that auction. The final price was determined by the second-price rule, that is, the bidder who submitted the highest bid won and paid (at most) a small increment ($.25) above the highest bid of the opponent, or, if the opponent did not bid, the price was the minimum bid of $1.[6] 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 influence how late some bidders can bid.

---

[4] Rasmusen (2003) shows that in a private-value auction model with costs of estimating one's own value, late bidding arises as the result of a bidder's incentive to avoid stimulating other bidders to examine their values. Hossain (2004) studies a private-value eBay auction model with informed bidders, who know their private valuation, and uninformed bidders, who know only whether or not their valuation exceeds the current price. In this model, sniping may be a best response of the informed bidders to the "learning by bidding" strategy of the uninformed bidders. In our experimental environment, however, all subjects are told their values at no cost before the auction starts.

[5] Others have noted deadline effects in Internet auctions (see, among others, Bajari and Hortaçsu, 2003; Malhotra and Murnighan, 2000; and Wilcox, 2000), and similar deadline effects have been noted in studies of bargaining (see, among others, Gächter and Riedl, 2005; Güth, Levati, and Maciejovsky, 2005; Roth, Murnighan, and Schoumaker, 1988). Gjerstad (2003) and Gjerstad and Dickhaut (1998) study the timing of bids and asks and the consequences for market outcomes in double auctions.

[6] As in Internet auctions, 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 submit the highest bid, the bidder who submitted his bid first is the high bidder at a price equal to the tie bid. 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 bid without raising the price.

Because eBay and Amazon are online auctions, it would have been possible to run 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. (2004). 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.

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 (typically the minimum increment over second-highest bid) are displayed to all. Stage 1 ends only after a period during which no player makes a bid. This design feature ensures that there is always time to respond to a bid submitted "early" in the auction, as it is the case in the theoretical models outlined in Ockenfels and Roth (forthcoming) and in Appendix B.

Stage 2 of the eBay.8 auctions consists of a single period. The bidders have the opportunity to submit one last bid with a probability $p = .8$ 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 in eBay.8.

Amazon. 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 $p = .8$. However, a successfully submitted stage-2 bid starts stage-1 bidding again (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.[7] Also, each bid had to meet or exceed the current minimum acceptable bid, which was $1 if no bid has been submitted previously, or the smallest increment ($.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 information policies employed by Amazon and eBay 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 1). eBay1, on the other hand, allows us to study the impact of this risk and therefore separate different contributory causes of late bidding (as we will explain below). Third, successfully submitted late bids 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, the second-price rule allowed a bidder in the experiments to bid by proxy. That is, a bidder could submit a bid early in the auction and have the resulting price register as the minimum

---

[7] 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 bid equals the current price.

TABLE 1   Experimental Treatments

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%

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 1 summarizes our experimental auctions.

☐   **Experimental procedure.** The study was conducted with 30 groups of 6 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 caused some auctions to have three bidders and others to have one, which rendered the data in all eBay.8 auctions numbers 19 and 20 incomparable. 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 with the sessions in which no problems occurred.

Auctions were run on networked computers using the z-Tree software toolbox by Fischbacher (1998). In each treatment, all rules were publicly explained with the help of example auctions (the instructions can be found in Ariely, Ockenfels, and Roth, 2004). In each auction, 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 the second-highest submitted bid). Participants were paid their cumulative earnings in all the auctions in which they participated, plus a show-up fee of $5 plus an additional $5 if they were at least five minutes early.
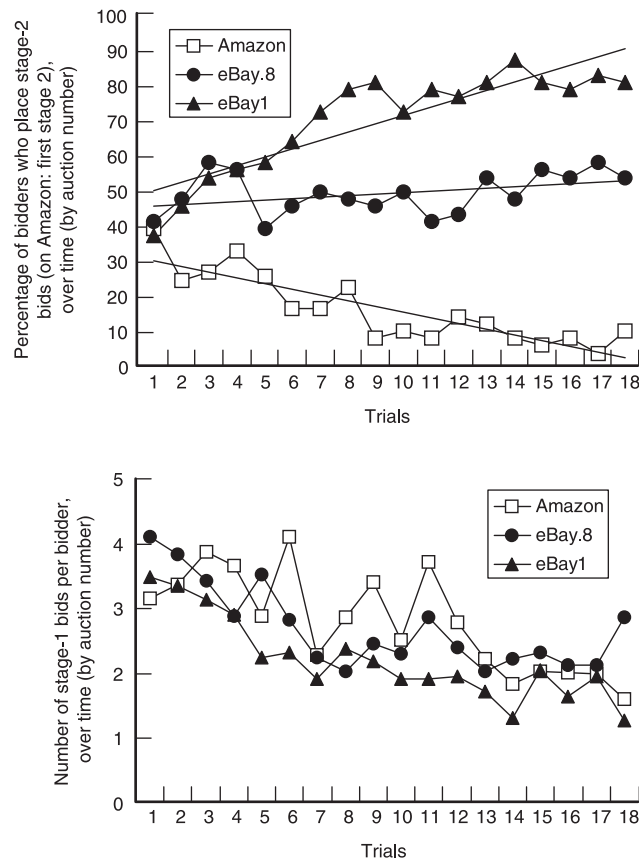
3. Experimental results

■   **The evolution of late and early bidding.** Figure 1 shows that the experimental results reproduce the main Internet observation: there is more late bidding in the fixed-deadline (eBay) conditions than in the automatic extension (Amazon) condition, and, 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.

The first panel of Figure 1 graphs the amount of late bidding, or sniping, by recording the percentage of bidders who place a bid in stage 2.[8] Since the measures used in Figure 1 are limited to at most one stage-2 bid 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 multiperiod auction conditions starts with about 40% of bidders submitting stage-2 bids, but by trial 18 Amazon has only about 10%, eBay.8 has 50%, and eBay1 has 80% 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

---

[8] In Amazon, there may be multiple stage-2 periods within an auction. In Figure 1 we included only the first stage 2 that determines whether the auction is extended at least once. Figure 1 includes stage-2 bids that were lost, which happened with probability .2 on eBay.8 and Amazon. The first panel includes lines that show the results of simple OLS regressions.

FIGURE 1

NUMBER OF BIDS PER BIDDER AND AUCTION OVER TIME

population (a Kruskal-Wallis H-test based on the 8 independent sessions for each auction type yields $p < .001$). Overall, there are weakly significantly more stage-2 bids on eBay1 than on eBay.8 (two-sided Mann-Whitney U-test, $p = .058$), and there are significantly more stage-2 bids in each of the eBay auction types than on Amazon ($p < .001$, for each comparison separately).[9]

The second panel of Figure 1 graphs the number of stage-1 bids per bidder over time. Comparison of the two panels 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 multiperiod auctions.[10]

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 number of bids between the auction conditions (Kruskal-Wallis H-test, $p = .125$).

Figure 2 shows that stage-1 bids are rarely placed by the current high bidder; early bids

---

[9] Restricting the analysis to experienced bidders (trials 10–18), the difference between eBay.8 and eBay1 becomes significant at the 1% level. A probit analysis in Table A1 in Appendix A confirms the time trends are highly significant: in both eBay conditions, the trend is toward more late bidding as bidders gain experience, while on Amazon, experienced bidders submit fewer late bids. Similarly, simple OLS regressions to estimate the time trend based on session-level observations reveal highly significant differences between slope coefficients.

[10] A closer look at the data 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 (forthcoming) 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 HIGHER BIDDER, AVERAGE ACROSS ALL
PERIODS AND TRIALS

are made mostly 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.[11]

Our observations with respect to the bid timing not only replicate the field observations but also reflect the underlying game-theoretic incentives. Ockenfels and Roth (forthcoming) argue that there are multiple reasons why sniping may be a rational strategy (even) in a private-value environment. 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 B we elaborate on this idea in a simplified model of implicit collusion. 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 B).

Sniping may also be a best response to incremental bidding that is observed both in the field (see Ockenfels and Roth, 2002 and forthcoming) 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. 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 interdependent valuations (Bajari and Hortaçsu, 2003), or perhaps they can learn about their private valuations by bidding incrementally (Hossain, 2004). Alternatively, increased attachment (such as the endowment effect) or competitive arousal can yield higher valuations over time (Heyman, Orhun, and Ariely, forthcoming; Ku, Malhotra, and Murnighan, 2004; Ockenfels and Ortmann, 2004). None of these explanations apply to our experimental environment, because values are exogenously induced and independent. Incremental bidding might also be caused by naive, 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 (forthcoming) as well as the lab evidence in Figure 1 suggests that multiple bidding is negatively correlated with experience. Bidding late on eBay may be a best reply to incremental bidding, because this strategy would not

[11] While there is no difference across auction types with respect to stage 1 (Kruskal-Wallis H-test, $p = .977$), there are significant differences with respect to stage 2 ($p < .001$). In particular, there are more high bidders submitting stage-2 bids in both eBay types, respectively, than on Amazon (two-sided Mann-Whitney U-test, $p = .001$ for each comparison separately), and there are more snipes by high bidders on eBay1 than on eBay.8 ($p = .045$).

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 B for a more formal argument).

Our data support the view that, in our eBay condition, early bidding does not pay: a bidder's payoff is significantly negatively correlated with his own number of stage-1 bids (the Spearman rank correlation coefficient is $-.172$, $p < .001$, for eBay1 and $-.113$, $p < .014$, for eBay.8), while the corresponding coefficient for the Amazon condition is not significant. Moreover, 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 sniping is occurring primarily 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 B). The evidence shown in Figure 1 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 $p < 1$.

☐   **The size of late and early bid increments, and price discovery.** Averaging over all bids in each stage, the three panels of Figure 3 summarize how much bids exceed the current minimum bid required, which is either the current price plus the minimum increment of $.25, or, before any bid is submitted, the reservation price of $1. 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.
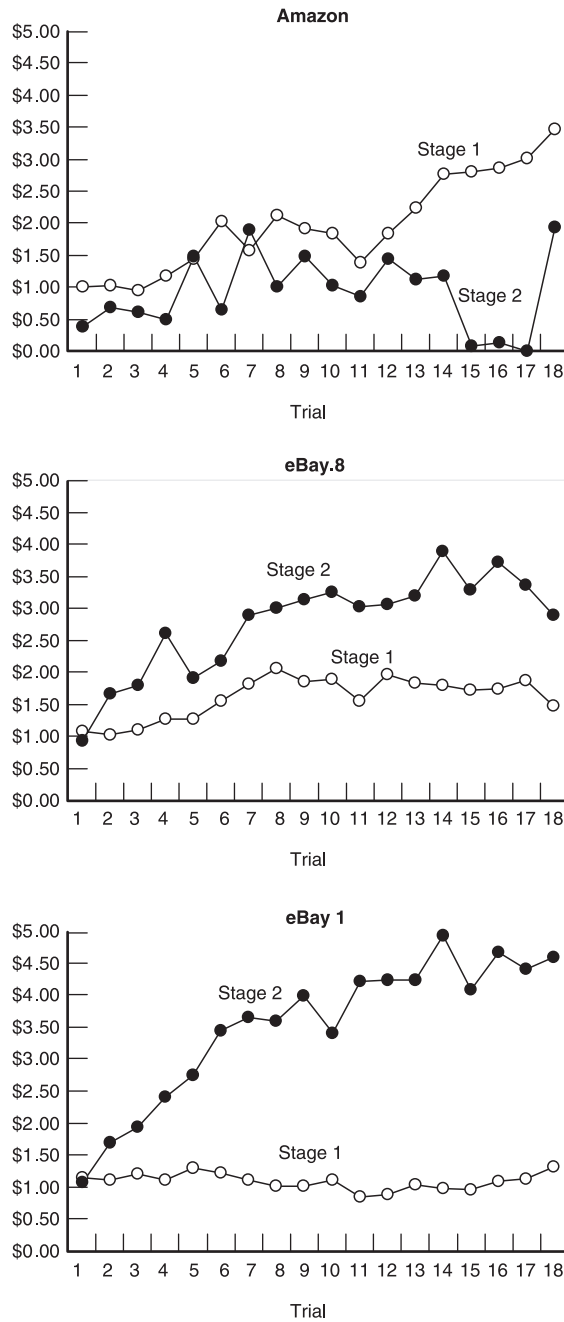
The first panel of Figure 3 shows that the average size of the bid increments placed in stage 1 on Amazon clearly grows over time, while the size of bid increments in stage 2 does not reveal a strong time trend. This is consistent with our earlier observation that the *numbers* of both stage-1 and stage-2 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 the first panel), they bid in larger increments in stage 1.

The situation is almost the opposite in each of the two eBay conditions. The second and third panels of Figure 3 show that although the average stage-1 bid increment stays relatively constant over time (slightly increasing on eBay.8 and about constant on eBay1), the average stage-2 bid increment strongly grows in both eBay conditions.[12] Moreover, these measures 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 stage-2 bids are becoming more frequent and stage-1 bids 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 eight 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 ($p = .001$ for each comparison separately), and significantly larger on eBay1 than on eBay.8 ($p = .021$). Also, although on Amazon stage-1 increases are significantly larger than stage-2 increases ($p = .012$), the opposite is true in each of the eBay conditions ($p = .059$ and $.002$ 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

---

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

FIGURE 3

AVERAGE INCREASE OF BIDS (CONDITIONED ON BIDDING) OVER CURRENT MINIMUM
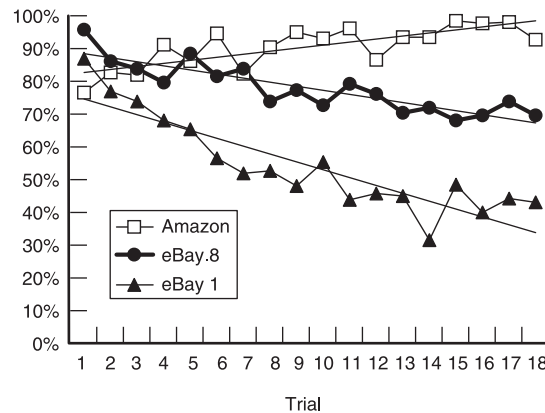BID

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.

    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

FIGURE 4

FINAL STAGE-1 PRICE (ON AMAZON: FIRST STAGE 1) AS PERCENTAGE OF FINAL PRICE
AND LINEAR TRENDS

reached more than 90% of the final price), whereas the opposite is true on eBay.8 (about 70%)
and eBay1 (less than 50%).[13]

☐   **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. 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
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.[14] However, since a winner can never advantageously influence
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 open 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.[15] More important, a bidder on Amazon who is cur-
rently 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.

Figure 5 shows that in all treatments the median of final bids relative to values is increasing
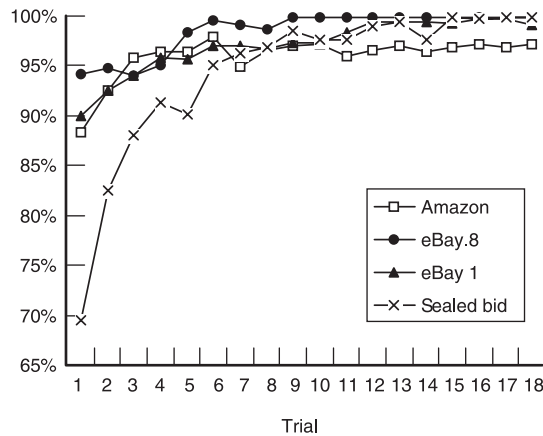
---

[13] The Spearman rank correlation coefficient between final (first) stage-1 price and final price is highest on Amazon
(.947), lower on eBay.8 (.570), and lowest on eBay1 (.270). All correlations are significant at $p = .001$.

[14] Recall that the rules determine that the price can never exceed the highest submitted bid.

[15] 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 $1.20, but to submit $100 immediately after
the price exceeded $1.20. 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, $1.10 in period 1, yielding a final price of $1.10.

FIGURE 5

MEDIAN OF FINAL BIDS (INCLUDING LOST STATE-2 BIDS) AS A PERCENTAGE OF VALUE

over time and, as predicted, never exceeds 100%.[16] 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% of values.[17] On Amazon, on the other hand, the median bid of experienced bidders stays below 100%. This is consistent 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 leads 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 (This is the reason why sniping is not a best response to incremental bidding on Amazon). 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.
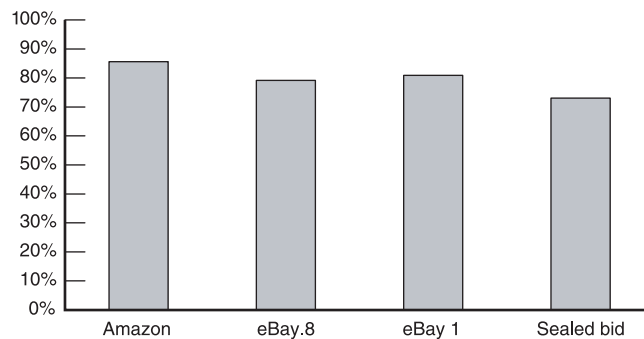
☐ **Revenue and efficiency.** Figure 6 shows the efficiency across all conditions measured as the average frequency the auctions are won by the bidder with the higher value, and Figure 7 shows median revenues.

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% 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% level). This seems to reflect the fact that Amazon is the only treatment in which low

---

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

[17] The average proportion of bidders whose final bid is equal to the private value is lowest in the Amazon (.125 and .391 if we include all final bids that are within a 25-cent range of the value) and sealed-bid (.151 and .366) conditions, and higher on eBay1 (.234 and .490) and eBay.8 (.251 and .586). In trial 18 the corresponding numbers are .167 and .396 for Amazon, .111 and .528 for sealed bid, .313 and .625 for eBay1, and finally .333 and .729 for eBay.8. Compared to previous second-price auction experiments (e.g., Kagel, Harstad, and Levin, 1987; Kagel and Levin, 1993), we observe rather modest overbidding (bidding more than one's private value). In particular, we have less than 10% overbidding in each of our open auction conditions and less than 20% overbidding in our sealed bid condition.

902   /   THE RAND JOURNAL OF ECONOMICS

FIGURE 6

AVERAGE EFFICIENCY

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.

The efficiency and revenue patterns across Amazon and eBay.8 are consistent with our observation from the subsection above on late and early bidding that the interaction between "naive" incremental and "sophisticated" bidders (who play best response against incremental bidding) influences the results. Incremental bidders may win Amazon auctions with bid amounts that are substantially below their values. At the same time, as mentioned above, on Amazon there is no pressure on incremental bidders to learn to bid up to their values, so they can safely stop bidding when they are the high bidder (since they can resume bidding if they are outbid). 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 naive 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 1 and 5). But since more bids are lost on eBay.8 than on Amazon (see Figure 1), both efficiency and revenue can be expected to be lower.[18]

FIGURE 7

MEDIAN REVENUES

[18] The median of the revenue relative to the second-highest value is 100.00% on Amazon and eBay.8, 99.37% on eBay1, and 97.23% in the sealed-bid condition.

Figures 6 and 7 illustrate that, aggregating over all trials and periods, efficiency and revenues in the two eBay conditions are statistically indistinguishable. This is inconsistent with our theory—the prediction that the risk that some bids will be suppressed in stage 2 in the eBay.8 condition should reduce efficiency and revenues. However, if we restrict ourselves to the analysis of experienced bidders (i.e., bidders in rounds 10–18), who snipe most often, the prediction is supported. In periods 10–18, average efficiency is 98% on eBay1 and 88% on eBay.8, and median revenues are $6.73 on eBay1 and $6.62 on eBay.8.

4. 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 that is reduced but not eliminated with experience, as was observed in the field data with a bidder's "feedback number" as an (imperfect) proxy for experience.

Our evidence is consistent with game-theoretic analyses that predict more late bidding in our eBay environment than in our Amazon treatment. While the experiment was not designed to exactly quantify the contributions of each of the multiple strategic reasons to snipe suggested by the theory, the fact that when the riskiness of sniping is removed (in the eBay1 treatment) the amount of late bidding goes up suggests that late bidding as a best response to incremental bidding is strongly present.

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 and the learning opportunities are better than those within the eBay and the sealed-bid conditions.[19]

Note that despite the superior control 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.

That said, it should be emphasized that great caution must be exercised about generalizing our results with respect to revenue and efficiency to the natural environment on the Internet. First, in order to focus on the difference in ending rules, our experiment controls away many differences (such as numbers of bidders, etc.) between eBay and Amazon that are important for determining the market outcomes. And second, efficiency and revenue depend on the risk of sniping that is set to 20% in our experiment. When this probability is smaller, or when the number of bidders is larger, the expected efficiency loss from sniping will be smaller, *ceteris paribus*.

Although 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 some complicating strategic factors as well as sources of variation across Internet auction sites such as endogenous and differential numbers of bidders, multiple

---

[19] We are not aware of earlier experiments conducted to test the performance of different online auction ending rules. However, some researchers experimentally studied the efficiency of English versus second-price auctions, including Coppinger, Smith, and Titus (1980), Cox, Roberson, and Smith (1982), Kagel, Harstad, and Levin (1987), and Kagel and Levin (2001). Lucking-Reiley (1999) compares revenues in second-price and English auctions in a field experiment but cannot compare efficiency because of a lack of information about individual values. Katok and Roth (2004) study the efficiency of different Dutch auction rules for multiple unit auctions.

items offered simultaneously, information externalities arising from affiliated values or uncertain (private) values, and heterogeneity of sellers, bidders, and products. By eliminating these factors, the experiment showed that they are neither necessary to produce sniping on eBay nor to produce the observed differences between eBay and Amazon:[20] the rules for ending these auctions are sufficient to drive the bidding dynamics. However, some of the factors that we eliminated from the experiment could well 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.

All field studies that we are aware of find substantial late bidding in hard-close auctions, but they typically differ in their explanations of sniping. Bajari and Hortaçsu (2003) explain sniping with a common-value auction model. Schindler (2003) finds that the simultaneity of auctions with similar items may explain late bidding. Wang (2003) shows theoretically that sequentially played identical eBay auctions create incentives to bid late and offer some field evidence. Hossain (2004) explores a model in which sniping is a best response to "learning-by-bidding" strategies performed by bidders who bid incrementally because they are not completely aware of their private valuation. Ku, Malhotra, and Murnighan (2004) explain their field data with the help of a model of emotional decision making and competitive arousal. Ockenfels and Roth (forthcoming) add some more potential reasons such as protection against "shill bidding" (an illegal attempt by sellers to raise the price by bidding just below the highest proxy bid) to this list. Hasker, Gonzales, and Sickles (2001) present statistical work on what they call "snipe or war" strategy. Analyzing eBay field data based on a theoretical private-value model adapted from Ockenfels and Roth (forthcoming), they find that sniping behavior appears not to be at equilibrium. However, they also find substantial amounts of multiple bidding. This suggests that, as in our experimental results, a large part of sniping behavior may be a response to (out-of-equilibrium) incremental bidders. We have focused here on the rules for ending the auction precisely because they affect the incentives for late bidding that may arise from many different causes.

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.[21] In short, the evidence suggests that there are multiple causes of late bidding in auctions, and that the strategic incentives to bid late are very much amplified by a hard close, i.e., by a rule that ends the auction at a fixed time regardless of bidding activity.

Appendix A

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

---

[20] See Bajari and Hortaçsu (2004) for a survey of the research dealing with some of these factors and other economic insights from Internet auctions.

[21] See, e.g., Roth and Xing (1994), Niederle and Roth (2003), and Fréchette, Roth, and Ünver (2004) for the effect of rules on timing and efficiency in other kinds of markets than auctions.

TABLE A1   **Probit Regression: Stage-2 Bids**

Independent Variables	Coefficients
Constant[a]	−.287**
	(−2.927)
Trial number (between 1 and 18) if Amazon and 0 otherwise[a]	−.101**
	(−12.249)
Trial number if eBay.8 and 0 otherwise	.020**
	(3.575)
Trial number if eBay1 and 0 otherwise	.104**
	(13.996)
ρ[b]	.486**
(Random effects)	(13.161)
Number of observations[c]	2,592
Log-likelihood	−1,184.896

Note: Random-effects probit model. Maximum-likelihood estimates (and *t*-statistics). Dependent variable = 1 for stage-2 bid (on Amazon: first stage 2) per bidder and per auction, and 0 otherwise. All three pairwise comparisons of the effects of trial numbers across treatments yield significance at the 1% level.

*Denotes significance at the 5% level (two sided), **Significance at the 1% level (two sided).

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

[b] Individual subject differences are clearly present as indicated by the highly significant ρ, the Hausman test statistic for the presence of random effects.

[c] Each bidder in each auction is one observation, making a total of 2,592 observations (= 3 treatments ∗ 48 bidders per treatment ∗ 18 trials).

Appendix B

■   **Theoretical considerations in a simplified model.** In this Appendix we simplify the auction environment 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. In their models, Ockenfels and Roth (forthcoming) take the price increment into account but restrict themselves to a particular distribution of values.

☐   **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.[22] 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 (forthcoming),[24] consider the following late-bidding strategies, which we will show constitute an equilibrium for risk-neutral bidders in 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 price war at which the equilibrium calls for a player to respond by bidding his true value in the subsequent period.

---

[22] 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 (Roth and Ockenfels, 2002).

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

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 $p = .8$ 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 .16 (= $p(1 − p)$), and he earns $4 (= value − opponent's value) if both stage-2 bids are successfully transmitted, which happens with probability .64 (= pp). If bidder 1's bid is lost at stage 2, which happens with probability .2 (= $1 − p$), his payoff is zero. 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 .16, and zero otherwise. Overall, then, we get a $4 expected payoff for bidder 1, and an $.80 expected payoff for bidder 2.

Unilateral deviation from the sniping strategy is not profitable for either bidder. First, in stage 2, any bid other 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 $.16 * (v_1 − \$1) + .64 * (v_1 − v_2)$, 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 (= $−.2v_1 + .36v_2 − .16$) 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. *Q.E.D.*

☐ **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).[24]

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 (forthcoming).[25] At a perfect Bayesian equilibrium in undominated strategies:

(i) 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_t at period t.

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

(iii) 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 be reached only 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. Since 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.

---

[24] 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 opportunity to warn against overinterpreting 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.

[25] As in the eBay case, their example is characterized by identical private values.

(iv) *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.   *Q.E.D.*

☐   **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 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 (naive) 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 \$1.25, 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 .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.   *Q.E.D.*

The "pure" form of incremental bidding as assumed for the proposition 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, the mechanics of the proof are robust to 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.

References

ARIELY, D., OCKENFELS, A., AND ROTH, A.E. "An Experimental Analysis of Ending Rules in Internet Auctions." Working Paper, Department of Economics, University of Cologne, 2004 (available at: http://ockenfels.uni-koeln.de/download/papers/aor.pdf).

ASKER, J., GROSSKOPF, B., MCKINNEY, C.N., NIEDERLE, M., ROTH, A.E., AND WEIZSÄCKER, G. "Teaching Auction Strategy Using Experiments Administered via the Internet." *Journal of Economic Education*, Vol. 35 (2004), pp. 330–342.

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

——— AND ———. "Economic Insights from Internet Auctions." *Journal of Economic Literature*, Vol. 42 (2004), pp. 457–486.

BOLTON, G., KATOK, E., AND OCKENFELS, A. "How Effective Are Online Reputation Mechanisms? An Experimental Investigation." *Management Science*, Vol. 50 (2004), pp. 1587–1602.

COPPINGER, V.M., SMITH, V.L., AND TITUS, J.A. "Incentives and Behavior in English, Dutch and Sealed-Bid Auctions." *Economic Inquiry*, Vol. 18 (1980), pp. 1–22.

COX, J.C., ROBERSON, B., AND SMITH, V.L. "Theory and Behavior of Single Object Auctions." In V.L. Smith, ed., *Research in Experimental Economics*. Greenwich, Conn.: JAI Press, 1982.

DELLAROCAS, C. "Reputation Mechanisms." In T. Hendershott, ed., *Handbook on Economics and Infdormation Systems), forthcoming*.

FISCHBACHER, U. "z-Tree—Zurich Toolbox for Readymade Economic Experiments—Experimenter's Manual." Working Paper no. 21, University of Zurich, 1999.

908 / THE RAND JOURNAL OF ECONOMICS

FRÉCHETTE, G., ROTH, A.E., AND ÜNVER, M.U. "Unraveling Yields Inefficient Matchings: Evidence from Post-Season College Football Bowls." Working Paper, Harvard University, 2004.

GÄCHTER, S. AND RIEDL, A. "Moral Property Rights in Bargaining with Infeasible Claims." *Management Science*, Vol. 51 (2005), pp. 249–263.

GJERSTAD, S. "The Impact of Pace in Double Auction Bargaining." Working Paper, University of Arizona, 2003.

——— AND DICKHAUT, J. "Price Formation in Double Auctions." *Games and Economic Behavior*, Vol. 22 (1998), pp. 1–29.

GÜTH, W., LEVATI, M.V., AND MACIEJOVSKY, B. "Deadline Effects in Ultimatum Bargaining: An Experimental Study of Concession Sniping with Low or no Costs of Delay." *International Game Theory Review*, Vol. 7 (2005), pp. 117–135.

HASKER, K., GONZALES, R., AND SICKLES, R.C. "An Analysis of Strategic Behavior and Consumer Surplus in eBay Auctions." Working Paper, Rice University, 2001.

HEYMAN, J., ORHUN, Y., AND ARIELY, D. "Auction Fever: The Effect of Opponents and Quasi-Endowment on Product Valuations." *Journal of Interactive Marketing*, Vol. 18 (2004), pp. 7–21.

HOSSAIN, T. "Learning by Bidding." Working Paper, Economics Department, Princeton University, 2004.

KAGEL, J.H. AND LEVIN, D. "Independent Private Value Auctions: Bidder Behaviour in First-, Second-, and Third-Price Auctions with Varying Numbers of Bidders." *Economic Journal*, Vol. 103 (1993), pp. 868–879.

——— AND ———. "Behavior in Multi-Unit Demand Auctions: Experiments with Uniform Price and Dynamic Vickrey Auctions." *Econometrica*, Vol. 69 (2001), pp. 413–454.

———, HARSTAD, R.H., AND LEVIN, D. "Information Impact and Allocation Rules in Auctions with Affiliated Private Values: A Laboratory Study." *Econometrica*, Vol. 55 (1987), pp. 1275–1304.

KATOK, E. AND ROTH, A.E. "Auctions of Homogeneous Goods with Increasing Returns: Experimental Comparison of Alternative 'Dutch' Auctions." *Management Science*, Vol. 50 (2004), pp. 1044–1063.

KU, G., MALHOTRA, D., AND MURNIGHAN, J.K. "Competitive Arousal in Live and Internet Auctions." Working Paper, Kellogg School of Management, Northwestern University, 2004.

LUCKING-REILEY, D. "Using Field Experiments to Test Equivalence Between Auction Formats: Magic on the Internet." *American Economic Review*, Vol. 89 (1999). pp. 1063–1080.

MALHOTRA, D. AND MURNIGHAN, J.K. "Milked for All Their Worth: Competitive Arousal and Escalation in the Chicago Cow Auctions." Working Paper, Kellogg School of Management, Northwestern University, 2000.

MILGROM, P. *Putting Auction Theory to Work*. New York: Cambridge University Press, 2004.

NIEDERLE, M. AND ROTH, A.E. "Unraveling Reduces Mobility in a Labor Market: Gastroenterology With and Without a Centralized Match." *Journal of Political Economy*, Vol. 111 (2003), pp. 1342–1352.

OCKENFELS, A. "Bidding in Private Value eBay Auctions: A Controlled Field Experiment." Working Paper, Economics Department, University of Cologne, 2004.

——— AND ORTMANN, A. "Bidding in Private Value eBay Auctions: A Controlled Field Experiment." Working Paper, Economics Department, University of Cologne, 2004.

——— AND ROTH, A.E. "Late and Multiple Bidding in Second Price Internet Auctions: Theory and Evidence Concerning Different Rules for Ending an Auction." *Games and Economic Behavior*, forthcoming.

——— AND ———. "The Timing of Bids in Internet Auctions: Market Design, Bidder Behavior, and Artificial Agents." *Artificial Intelligence*, Vol. 23 (2002), pp. 79–87.

RASMUSEN, E. "Strategic Implications of Uncertainty Over One's Own Private Value in Auctions." Working Paper, Indiana University, 2003.

RESNICK, P. AND ZECKHAUSER, R. "Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay's Reputation System." In M.R. Baye, ed., *The Economics of the Internet and E-Commerce. Advances in Applied Microeconomics*, Vol. 11. Boston: JAI, 2002.

ROTH, A.E. "The Economist as Engineer: Game Theory, Experimental Economics and Computation as Tools of Design Economics." *Econometrica*, Vol. 70 (2002), pp. 1341–1378.

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

——— AND XING, X. "Jumping the Gun: Imperfections and Institutions Related to the Timing of Market Transactions." *American Economic Review*, Vol. 84 (1994), pp. 992–1044.

———, MURNIGHAN, J.K., AND SCHOUMAKER, F. "The Deadline Effect in Bargaining: Some Experimental Evidence." *American Economic Review*, Vol. 78 (1988), pp. 806–823.

SCHINDLER, J. "Late Bidding on the Internet." Working Paper, University of Vienna, 2003.

WANG, T.J. "Is Last Minute Bidding Bad?" Working Paper, UCLA, 2003.

WILCOX, R.T. "Experts and Amateurs: The Role of Experience in Internet Auctions." *Marketing Letters*, Vol. 11 (2000), pp. 363–374.

VICKREY, W. "Counterspeculation, Auctions, and Competitive Sealed Tenders." *Journal of Finance*, Vol. 16 (1961), pp. 8–37.