

Late and Multiple Bidding in Second Price Internet Auctions: Theory and Evidence Concerning Different Rules for Ending an Auction

Axel Ockenfels & Alvin E. Roth*

May 27, 2003

Abstract

In second price internet auctions with a fixed end time, such as those on eBay, many bidders ‘snipe’, i.e., they submit their bids in the closing minutes or seconds of an auction. Late bids of this sort are much less frequent in auctions that are automatically extended if a bid is submitted very late, as in auctions conducted on Amazon. We propose a model of second price internet auctions, in which very late bids have a positive probability of not being successfully submitted, and show that sniping in a fixed deadline auction can occur even at equilibrium in auctions with private values, as well as in auctions with uncertain, dependent values. Sniping in fixed-deadline auctions also arises out of equilibrium, as a best reply to incremental bidding. However, the strategic advantages of sniping are eliminated or severely attenuated in auctions that apply the automatic extension rule. The strategic differences in the auction rules are reflected in the field data. There is more sniping on eBay than on Amazon, and this difference grows with experience. We also study the incidence of multiple bidding, and its relation to late bidding. It appears that one substantial cause of late bidding is as a strategic response to incremental bidding. (*JEL* C73, C90, D44)

* Ockenfels: Max Planck Institute for Research into Economic Systems, Strategic Interaction Unit, Kahlaische Str. 10, D-07745 Jena, Germany; ockenfels@mpiew-jena.mpg.de (www.mpiew-jena.mpg.de/esi/ockenfels). Roth: Harvard University, Department of Economics and Graduate School of Business Administration, Soldiers Field Road, Baker Library 183, Boston, MA 02163, USA; aroth@hbs.edu (www.economics.harvard.edu/~aroth/alroth.html). We thank the Associate Editor of this special issue, Eric van Damme, and two anonymous referees for very helpful comments. This project received financial support from the National Science Foundation, the Harvard Business School, and the Deutsche Forschungsgemeinschaft through the Emmy Noether Program.

I. Introduction

Some of the most successful internet auction houses such as eBay and Amazon employ second price auctions, in which the high bidder wins the object being auctioned, but pays a price equal to a small increment above the second highest bid. A bidder may submit a reservation price, and this (maximum) price is used to bid for him by proxy. At any point in the auction the “current price” is a minimum increment above the second highest submitted reservation price. Hence, the bidder with the highest reservation price submitted during the auction will win the auction and pay only the minimum increment above the second highest submitted reservation price. One of two ending rules is generally applied. Some auction houses such as eBay have a fixed end time (a “hard close”), while others, such as Amazon, which operate under otherwise similar rules, are automatically extended if necessary past the scheduled end time until some number of minutes (ten on Amazon) have passed without a bid (a “soft close”).

Roth and Ockenfels (2002) noted that there is a great deal of very late bidding (called “sniping”) on eBay, and considerably less on Amazon. In the present paper, we investigate this phenomenon more closely, theoretically and empirically. We also investigate a distinct but, as we will argue, closely related phenomenon: many bidders submit multiple bids in the course of the auction, i.e., they submit and later raise the reservation price they authorize for proxy bidding on their behalf.

A number of observers have expressed surprise and puzzlement at these patterns of late and multiple bidding in internet auctions that arise despite advice from auctioneers that bidders should simply submit their maximum willingness to pay, once, early in the auction.¹ One source of the puzzlement at why the “single early bid” strategy is not predominant is the fact that in a second price sealed-bid private value auction, as first studied by Vickrey (1961), it is a *dominant* strategy for bidders to submit as their reservation price their maximum willingness to pay. In

¹ Landsburg (1999), for instance, describes his own multiple-bid and last-minute-bidding behavior and concludes: “Maybe eBay just makes me giddy. As a free market aficionado, I am intoxicated by the prospect of one-stop shopping for houses, cars, Beanie Babies, and underwear, all at prices that adjust instantly to the demands of consumers around the globe. Or maybe the behavior of eBayers can be explained only by subtler and more carefully tested theories that have not yet been devised.” And eBay “encourage[s] all members to use the proxy bidding system that is in place to bid the absolute maximum they are willing to pay for an item right from the start and let the proxy bidding system work for them.” (http://webhelp.ebay.com/cgi-bin/eHNC/showdoc-ebay.tcl?docid=88269&queryid=proxy_bidding).

this view, bidding less than one's true value in a private value auction is an error, and the observed behavior might primarily be due to naïve, inexperienced, or plain irrational behavior (see, e.g., Wilcox, 2000, for this view of eBay auctions, and Malhotra and Murnighan, 2002, who also conclude they are seeing irrational late bidding, in an interesting study of a Chicago charity auction for artists' renderings of cows).

Sniping also involves a non-negligible risk. Because the time it takes to place a bid may vary considerably due to erratic internet traffic or connection times, last-minute bids have a positive probability of being lost. In a survey of 73 bidders who successfully bid at least once in the last minute of an eBay auction, 63 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 (Roth and Ockenfels, 2002). Human and artificial bidders do not differ in this respect. The last-minute bidding service esnipe.com, which offers to automatically place a predetermined bid a few seconds before the end of an eBay auction, acknowledges that it cannot make sure that all bids are actually received on time by eBay (<http://www.esnipe.com/faq.asp>, 2000).

The present paper will show that the incentives for late bidding are nevertheless quite robust in fixed-deadline auctions. Not only is late bidding a best response when there are naïve, incremental bidders, but late bidding may be consistent with equilibrium, perfectly rational behavior both in auctions with private values, and auctions with uncertain, dependent values. However, we will also show that rational late bidding is sensitive to the rules of how the auction ends. This will motivate the empirical part of the study.

We analyze data collected from completed eBay and Amazon auctions. Auctions conducted by Amazon differ from those conducted by eBay in that, if there is a bid in what would have been the last ten minutes of the auction, then the auction is automatically extended so that it does not end until ten minutes after the final bid. This removes the strategic advantage of bidding at the last minute, so that the late-bidding strategies that are in equilibrium or a best response to naïve behavior on eBay will not be on Amazon. The data on the timing of bids reflect this strategic difference in various ways. Not only is there much more late bidding on eBay than on Amazon, but more experienced bidders on eBay submit late bids more often than do less experienced bidders, while the effect of experience on Amazon goes in the opposite direction. Further, while there is late bidding both by bidders who bid only once and by those who bid

multiple times, bidders who bid only once bid later, and are more experienced than bidders who bid multiple times.

II. Last minute bidding in second price internet auctions with a fixed deadline

II.1 A strategic model of the eBay bidding environment

We consider here the standard eBay auction format, with a specified minimum opening bid but no (other) seller's reservation price.² We will construct a model with a multiplicity of equilibria, including equilibria with last-minute bidding, even when there is no uncertainty about own values and last minute bids run the risk of being lost. For clarity, we first present the strategic structure of the auction, to which we will add the players' valuations (private values or uncertain, dependent values) when we examine particular equilibria.

The strategic model (the eBay game-form):

- There are n bidders, $N = \{1, \dots, n\}$.
- There is a minimum initial bid m (the seller's reservation price), and a smallest increment s by which subsequent bids must be raised. We assume for simplicity that s is constant. (In fact, on eBay, the smallest increment increases as the price increases).
- The "current price" (or "high bid") in an auction with at least two bidders generally equals the minimum increment over the second highest submitted bidders' reservation price. There is one exception to this: The price never exceeds the highest submitted reservation price; if the difference between the highest and the second highest submitted reservation price is smaller than the minimum increment, the current price equals the highest reservation price.

² Sellers typically can choose to set not only a (public) minimum initial bid, but also an additional (secret) reservation price below which they will not sell. We will concentrate here on auctions in which the seller's reservation price is known, since bidding in an auction with an unknown reservation price presents additional strategic prospects. In particular, auctions in which the high bid is less than the reserve price sometimes lead to post-sale negotiations. Therefore, there is an incentive for an agent not to bid against himself if he is the current high bidder for fear of going over the reservation price and having to pay it in full.

- Each submitted reservation price must exceed both the current price and the corresponding bidder's last submitted reservation price (i.e., a bidder cannot lower his own previous reservation price).³
- The bid history lists the current price and the submission time of each bidder's last submitted reservation price⁴ along with the corresponding bidder ID $j \in N$. The highest reservation price is not revealed.
- A player can bid at any time $t \in [0, 1) \cup \{1\}$. A player has time to react before the end of the auction to another player's bid at time $t' < 1$, but the reaction cannot be instantaneous, it must be strictly after time t' , at an earliest time t_n , such that $t' < t_n < 1$. For specificity, this earliest reaction time is the first $t_n > t'$, chosen from a countably infinite subset $\{t_n\}$ of $[0, 1)$, such that the $\{t_n\}$ converge to 1. The information sets of the game are such that if t is between t_{n-1} and t_n ($t_{n-1} < t < t_n$), then the players know the bid histories up to t_{n-1} .⁵ That is, in the early part of the bidding, bidders can make their behavior in that half-open interval contingent on other bids observed in the interval— they have time to react to one another.
- If at some t more than one bidder submitted the highest reservation price, the bidder who submitted her bid first is the high bidder at a price equal to the reservation price. If multiple

³ If the current high bidder submits a new, higher reservation price, the current price is not raised, although the number of bids is incremented. The one exception is if the high bidder's reservation price was less than the minimum increment above the second highest. In this case, if the high bidder submits a new bid, it will raise the current price to the minimum increment above the second highest reservation price.

⁴ After our first data set (see Section IV) was collected, eBay changed this rule. Now eBay's bid history for each auction includes *all* bids. These changes gave us the opportunity to analyze the multiple bidding phenomenon in greater detail (see Section V), while at the same time they do not alter our theoretical conclusions. Amazon continues to include only each bidder's last bid.

⁵ This allows us to be sure that the path of play will be well defined for any strategy, and avoid the potential problem with continuous time models of having the path of play undefined, as when each player's strategy calls for him to react to the other players 'immediately.' A referee noted that one could also start from the outset with a countable number of decision times, like $\{1 - 1/n | n = 1, 2, \dots\} \cup \{1\}$, without any damage to our theoretical results. We prefer, however, the continuous time model as a somewhat more natural description of eBay. Also, in all following analyses we restrict ourselves to equilibria in undominated strategies. Because newly submitted reservation prices must rise at least by the minimum increment and because submitted reservation prices above value are dominated, this implies that the number of reservation prices submitted by each bidder is finite. However, the model must specify whether bidders are allowed to submit an infinite number of reservation prices, and if so how the outcome of the game is determined. There is an implicit rule on eBay (and similarly also on Amazon) that takes care of this problem: eBay's interface does not allow submitting reservation prices above \$9,999,999,999.99. This maximum bid size effectively limits the number of bids that can potentially be placed, and at the same time does not affect our theoretical results if we assume that the maximum value is below this maximum bid size.

bids are submitted simultaneously at the same instant t , then they are randomly ordered, and each has equal probability of being received first.

- At $t = 1$, everyone knows the bid history prior to t , and has time to make exactly one more bid, without knowing what other last minute bids are being placed simultaneously.
- Bids submitted before time $t = 1$ are successfully transmitted with certainty.
- At time $t = 1$, the probability that a bid is successfully transmitted is $p < 1$. The probability p is assumed to be exogenous, i.e., not influenced by the bidding in a given auction.⁶

II.2 Some theoretical observations about eBay auctions with private values

In this section we assume that each bidder j has a true willingness to pay, v_j , distributed according to some known, bounded distribution. A bidder who wins the auction at price h earns $v_j - h$, a bidder who does not win earns 0. At $t = 0$ each player j knows her own v_j .⁷ At every $t \in (0, 1) \cup \{1\}$, each player knows also the bid histories for all $t_n < t$.

Unlike the standard case of second price sealed bid auctions (as analyzed by Vickrey, 1961), there are generally no dominant strategies in the dynamic eBay environment.

Theorem [No dominant strategies]. A bidder in the second price eBay-model auction with private values does not have any dominant strategies (so long as his value is greater than $m + s$).

Proof. Let $n = 2$ (if $n > 2$ assume that all additional bidders do not bid at any time). It is sufficient to show that there is no strategy for bidder j who has value $v_j > m + s$ that is a best reply to every strategy of the other bidder, i .

⁶ It is obviously an approximation both to take a discontinuous drop from $p = 1$ to $p < 1$, and to model the “last minute” as a single point in time. Here we sacrifice strict realism for simplicity and clarity. Also, major server attacks may affect eBay’s server speed and thus the probability p . However, since according to eBay (see http://www.forbes.com/asap/2000/1127/134_2.html) there are on average about 1,200 bids per minute submitted to eBay, ‘regular’ bidding in any given auction can hardly have a measurable effect on the performance of eBay’s server.

⁷ Values may be independent (as in the standard private value auction environment) or dependent, but unlike the standard common value case bidders are assumed to know their values so that the maximum willingness to pay is not changed by information conveyed by others’ bids.

Suppose first that i 's strategy is to bid the minimum bid m at $t = 0$ and then not to bid further at any information set at which he remains the high bidder, but to bid B (with $B > v_j + s$) whenever he learns that he is not the high bidder. Against this strategy, j 's best reply is not to bid at any time $t < 1$, and to bid v_j at time $t = 1$. The payoff to j from this strategy is $p(v_j - m - s) > 0$, and it is easy to see that no other bid at time $t = 1$ could yield j a larger payoff.

But now suppose instead that player i 's strategy is not to bid at any time. Then a strategy for player j that calls for bidding only at $t = 1$ will give j the expected payoff $p(v_j - m) < v_j - m$, which is the payoff j could achieve from the strategy of bidding v_j at $t = 0$ (or at any $t < 1$). Hence no best reply to the strategy for player i considered in this paragraph, is among player j 's best replies to i 's strategy in the previous paragraph, and hence j does not have a dominant strategy. This completes the proof of the theorem.

Since the main driver of the above result is the dynamic aspect of the eBay game, one might wonder if in the truncated eBay-game that starts at $t = 1$ bidding one's true value is always a (weakly) dominant strategy because, as in Vickrey's sealed bid auction, at $t = 1$ the auction is over after bidders simultaneously submit their bids. Due to the minimum bid increment, however, this is not the case. Assume for instance that bidder i plans to bid some $b_i < v_j - s$ in the truncated game. In this situation bidder j would have to pay a price of $b_i + s$ if he bids his value and none of the bids is lost. But by bidding an amount slightly above b_i (and below $b_i + s$) he could avoid paying the entire minimum increment because the rules determine that the price can never exceed the highest submitted reservation price. Thus, bidder j wants to condition his bid on bidder i 's bid, implying that no dominant strategy exists in the truncated game that only starts at $t = 1$.⁸

⁸ Several remarks are in order. First, technically, if money is modeled continuously a best reply does not exist because bidder j wants to beat the highest bid by as small a margin as possible. Second, the truncated game will not in general be a game in standard extensive form, since the information sets at $t = 1$ are not singletons when players do not know each other's values (but the initial distribution over values can be applied to make it a standard game). Finally, note that since a winner can never advantageously affect the price in case of winning by bidding higher than his value, and since a winner can never push the price down by more than one increment by underbidding, bidding one's true value may be called an " s -dominant strategy" of the game starting at $t = 1$. That is, bidding one's true value is a strategy that always yields a payoff not more than the minimum increment s below the supremum achievable by any other strategy, regardless of the strategies chosen by the other bidders.

The theorem demonstrates that the eBay second price auction is very different from the textbook case of a sealed bid second price auction at which every player has a dominant strategy of bidding his maximum willingness to pay. Thus, the standard conclusions drawn from the analysis of second price auctions cannot be immediately applied to eBay bidding. In particular, since there are no dominant strategies in the eBay environment it is not surprising that bidding late may be a best response to all kinds of bidding behaviors. When we turn to the data (especially Section V), we will see that there is a good deal of incremental bidding. While the particulars of incremental bidding may vary, let us, for specificity, define incremental bidding as the strategy that calls for a bidder to bid m at $t = 0$ and then, whenever he is not the high bidder at some $t < 1$, to bid in minimum increments at some $t < 1$ until he (re)gains the high bidder status or until he reaches his value, whichever comes first.⁹ The following theorem shows that late bidding may be a best response to this kind of incremental bidding.

Theorem [Best response to incremental bidding on eBay]. Bidding at $t = 1$ can be a best reply to incremental bidding.

Sketch of proof. For simplicity, consider the case in which bidders' values are common knowledge. Suppose there are two bidders, one rational bidder j with value $v_j > m + s$, and one incremental bidder i with value $v_i > m + s$. Facing an incremental bidder, j faces a tradeoff. Bidding early makes sure that the bid is received, but induces a bidding war that raises the price. Bidding late gives the incremental bidder no time to push up the price (because he does not learn that he is no longer the high bidder until the game is over) but runs the risk that the bid and thus the auction is lost. However, whenever $v_i > v_j + s$ bidder j 's best reply is to bid v_j at $t = 1$ (other bids at $t = 1$ also qualify as best replies) and not to bid early, yielding j a payoff of $v_j - (m + s)$ with probability p . Any bid $b \leq v_j$ by j at some $t < 1$ would be subsequently outbid by i at some t' with $t < t' < 1$ (and any strategy that calls for bidder j to bid above value at t is dominated by the otherwise identical strategy in which j bids at most his value at t), and at the same time would

⁹ Reasons for incremental bidding include mistaken analogy with English first-price auction bidding (the data analyzed in Section V reveals that incremental bidders are relatively inexperienced), and psychological reasons that might cause a bidder to increase his maximum bid in the course of an auction because his maximum willingness to pay is increasing over time (see, e.g., Malhotra and Murnighan, 2002).

raise the best price that is achievable at $t = 1$ above $m + s$. (Whether late bidding is a best response or not in situations in which the rational bidder j 's value is above the incremental bidder i 's value will depend on the distribution from which values are drawn and on the parameters p , m and s .)

Even though much late bidding appears to be a strategic reaction to (non-equilibrium) incremental bidding,¹⁰ it is still of interest to know if late bidding can constitute equilibrium behavior despite the risk involved in late bidding. The next theorem shows that this is the case. This suggests that, even if bidders initially learn to bid late in response to unsophisticated behavior of others, late bidding may persist even as all bidders gain experience. The intuition behind last-minute bidding at equilibrium in auctions with private values will be that there is an incentive not to bid too high when there is still time for other bidders to react, to avoid a bidding war that will raise the expected final transaction price. And mutual delay until the last minute can raise the expected profit of all bidders, because of the positive probability that another bidder's last-minute bid will not be successfully transmitted. Thus at such an equilibrium, expected bidder profits will be higher (and seller revenue lower) than at the equilibrium at which everyone bids true values early.

Theorem [An example of equilibrium sniping in eBay auctions with private values]. There can exist perfect Bayesian equilibria in which bidders do not bid early, at some $t < 1$, but only bid at the last moment, $t = 1$, at which time there is only probability $p < 1$ that the bid will be transmitted.

Proof. Consider the case of two bidders, $N = \{1, 2\}$. For simplicity we assume that both bidders have the same value H (i.e. that values are drawn from the degenerate distribution with all its mass at H), with $H > m + 2s$.¹¹ Clearly, in such a situation, many equilibria exist. In particular,

¹⁰ See Section V for field evidence, Ariely et al. (2002) for experimental evidence, and Ockenfels and Roth (2002) for anecdotal evidence. Of course, late bidding may also be a best response to other incremental bidding (or 'price war') behaviors, including that of a dishonest seller who attempts to raise the price by using "shill bidders" to bid against a proxy bidder.

¹¹ This assumption is sufficient but not necessary for making sniping an equilibrium strategy in our model; see the discussion at the end of this section.

the “single early bid” strategy recommended by eBay corresponds to a set of equilibria in our environment. That is, strategies that call for submitting one’s value at some $t < 1$ (and bidding only once) are best responses to each other. In these equilibria, however, both bidders will earn zero. We will demonstrate that mutual sniping, and not bidding early, may constitute a perfect Bayesian Nash equilibrium (in fact, subgame perfect, since the distribution of values is degenerate) that gives the bidders positive profits from late bidding whenever the probability p that a bid at $t = 1$ is transmitted successfully is less than 1 and bigger than 0.

Consider the following bidding strategies, which we will show constitute an equilibrium for any probability $p > 0$ of a successfully transmitted bid at $t = 1$. On the equilibrium path, each bidder i ’s strategy is not to bid at any time $t < 1$, and to bid H at time $t = 1$, unless the other bidder deviates from this strategy. Off the equilibrium path, if player j places a bid at some $t' < 1$, then player i bids H at some $t > t'$ such that $t < 1$. That is, each player’s strategy is to do nothing until $t = 1$, unless the other bidder makes an early bid, that would start a price war at which the equilibrium calls for a player to respond by promptly bidding his true value. The threat of a price war is credible, because, as mentioned above, bidding values early also constitutes an equilibrium. In equilibrium, a bidder earns $H - m$ if his bid at $t = 1$ is successfully transmitted and the other bidder’s bid is lost, which happens with probability $p(1 - p)$. Otherwise, and whenever the bidder deviates from his equilibrium strategy by bidding at some $t < 1$, his payoff is zero, regardless of whether bids at $t = 1$ are lost or not. In the subgame starting at $t = 1$, and given that no bids were registered at $t < 1$, mutually bidding true values is a Nash equilibrium. This proves that the indicated strategies constitute a subgame perfect equilibrium (and thus a perfect Bayesian equilibrium).

The proof shows that even at equilibrium in auctions with private values, bidders may have strategic reasons to refrain from bidding as long as there is time for others to react, since otherwise they can cause a bidding war that raises the expected transaction price. In the proof we took the values of the bidders to be identical, but this is not essential. It does let the proof proceed without attention to the parameters of the problem. When bidders have non-identical values, late bidding at a perfect Bayesian equilibrium as in the example above is sensitive to the distribution of values, to the probability p , and to the number of bidders, all of which also affect

the expected profits at both early and late bidding equilibria (for a numerical example see Ariely et al., 2002). At the same time it is clear, however, that the theorem only shows that some things *can* happen in equilibrium, by way of example. We do not attempt to characterize the set of all possible Bayesian equilibria in this paper. Finally, because sniping involves the risk that bids are lost, both revenue and efficiency are affected. For instance, in the case of two bidders, identical values, and no multiple bidding in equilibrium considered above, an inefficient allocation results if all bids are lost (with probability $(1 - p)^2$), and the expected revenue is smaller than in a second price auction in which all bids are accepted with probability 1 if one or both bids at $t = 1$ are lost (with probability $1 - p^2$).

II.3 An example of eBay auctions with uncertain, dependent values

The simplest models of auctions with uncertain, dependent values treat the bidders as more or less symmetric, except for their private signals about the value, and sometimes for a private value component. Bidders can use the bids of the other bidders to update their information about the value of the good. In an English auction conducted so that bidders have to stay in the auction to remain active and can be observed to drop out, the dropout times of other bidders are informative. In an eBay auction, however, it cannot be observed if a bidder has dropped out; he might bid again, or for the first time, at $t = 1$. Bajari and Hortaçsu (forthcoming), for example, develop a common value model in which all bidders bid at $t = 1$, to avoid giving other bidders information. (In their model there is no cost to bidding at $t = 1$, however, since bids are transmitted with certainty.)

In general, late bids motivated by information about values arise either so that bidders can incorporate into their bids the information they have gathered from the earlier bids of others, or so bidders can avoid giving information to others through their own early bids. In eBay auctions, a sharp form of this latter cause of late bidding may arise when there is asymmetric information, and some players are better informed than others. The following example illustrates this kind of equilibrium late bidding. It is motivated by auctions of antiques, in which there may be bidders

who are dealers/experts who are better able to identify high value antiques, but who do not themselves have the highest willingness to pay for these once they are identified.¹²

In our simple “*dealer/expert model*”, the bidding structure is the same as above, with the seller’s reservation price represented by the minimum allowable initial bid, $m > 0$. The object for sale has one of two conditions, “Fake,” with probability p_F , or “Genuine,” with probability $p_G = 1 - p_F$. There are $n = 2$ (representative) bidders. The first, U , is uninformed but discerning, and values Genuines more highly than Fakes ($v_U(F) = 0 < v_U(G) = H$), but cannot distinguish them, i.e. cannot tell whether the state of the world is F or G . The second, I , is informed (e.g. an expert dealer), with perfect knowledge of the state of the world, and values $v_I(F) = 0$ and $v_I(G) = H - c$, with $0 < m < H - c + s < H$. Then the dependent value aspect of this auction arises from the concern of the bidders with whether the object is Genuine. The strategic problem facing the informed bidder is that if the object is Genuine and he reveals this by bidding at any time $t < 1$, then the uninformed bidder has an incentive to outbid him. But the strategic problem facing the uninformed buyer is that, if he bids H without knowing if the object is Genuine, he may find himself losing money by paying $m > 0$ for a Fake.

Theorem [An example of equilibrium sniping in eBay auctions with uncertain, dependent values]. When the probability that the object is Fake is sufficiently high, there is a perfect Bayesian equilibrium in undominated strategies in this “*dealer/expert model*,” in which the uninformed bidder U does not bid, and the informed bidder bids only if the object is Genuine, in which case he bids $v_I(G) = H - c$ at $t = 1$. If the informed player deviates and makes a positive bid at any $t < 1$, then the uninformed bidder bids H at some t' such that $t < t' < 1$.

Sketch of the proof. Any strategy that calls for the informed player to place a positive bid when the object is Fake is dominated by the strategy of placing no bid. Therefore, if the informed

¹² Some experienced antiques bidders in the survey by Roth and Ockenfels (2002) stated that they snipe to avoid information revelation. Of course, a bidder could also hide his identity behind different user names. In fact, eBay concedes that “Many of our users have determined that it is useful to have multiple user IDs for their businesses.” (<http://www.auctionwatch.com/awdaily/tipsandtactics/buy-multiple.html>). Since, on the other hand, changing identities comes at the cost of spreading the reputation across different identities, and maintaining different email addresses and credit cards, it is likely that identities are not always changed when information asymmetries would dictate so. Our evidence in Section IV supports this conjecture.

player ever makes a positive bid at some $t < 1$, the uninformed player can conclude that the object is genuine and has sufficient time to subsequently submit H at some t' such that $t < t' < 1$. As a consequence, in our equilibrium, the informed bidder never bids before $t = 1$. However, so long as the conditional expected value of the object, conditional on winning, is negative it never pays for the uninformed bidder to bid if the informed bidder has not already bid, since in this case the expected loss from winning a Fake and paying m is larger than the expected gain from winning a Genuine and paying m . Formally, given that the informed bidder does not bid before $t = 1$, the uninformed bidder will not submit a bid if

$$\frac{p_G(1-p)(H-m) - p_F m}{p_G(1-p) + p_F} < 0,$$

or equivalently if $p_F/p_G > (1-p)(H-m)/m$. Thus, if p_F is sufficiently large, no player will bid at $t < 1$, so that for the informed bidder bidding true value at $t = 1$ is part of a perfect Bayesian equilibrium strategy when the object is Genuine.

This shows that, in an eBay type auction, an (identifiable) informed bidder has an incentive to not make a high bid on a genuine item until the last minute. This incentive to snipe is robust with respect to the probability p (with $0 < p < 1$) that a last-minute bid is transmitted successfully, because the informed bidder cannot win by bidding early *regardless* of the value of p . Also, the condition on p_F/p_G is sufficient but not necessary for the informed bidder to snipe in equilibrium. Even if it is profitable for the uninformed bidder to bid early, i.e. if $p_F/p_G < (1-p)(H-m)/m$, or if $m = 0$ there can exist equilibria in which the uninformed bidder bids his (positive) conditional expected value early, at some $t < 1$, and the informed bidder, as before, bids only once, at $t = 1$, if the object is genuine.¹³ However, because there are many ways in which values can be dependent (as opposed to the “pure” or “mineral rights” model of common values), we do not seek here to develop a general model of dependent values. We simply re-emphasize that the demonstration in the previous section, that late bidding can occur at

¹³ If we allow weakly dominated strategies, there are furthermore equilibria at which the informed bidder bids, say, $\$1,000,000 + H$ at $t = 0$ when the object is genuine, and gets the object for m . Finally, we have treated here the case of a single auction. It is not a dominated strategy for a dealer bidding in many auctions to sometimes bid above his value.

equilibrium in auctions with private values, does not imply that late bidding is itself evidence of private values, any more than it is evidence of common values.

Note in particular that the “expertise effect” modeled here that causes experts to bid late can be additive with the “avoiding price war effect” already noted in the case with private values. That is, experts could bid late to avoid alerting non-experts of the value of the object, at the same time as they and other bidders bid late to avoid setting off price wars. Section IV will show that there is more late bidding in eBay-Antiques than in eBay-Computers. Since these categories might reasonably be expected to have a different scope for expert information, the expertise effect appears to be reflected in the timing of bids.

But first, we will use these simple examples to demonstrate how different the strategic incentives can be under the rules by which auctions are conducted on Amazon.

III. A simple model of bidding in internet auctions with the automatic extension rule

Based on the results in the last section, we now demonstrate by example that the value of sniping in auctions with a hard close may be eliminated in auctions with a soft close such as Amazon’s automatic extension rule.¹⁴ We caution from the start that our examples in this section do not imply that there is no combination of independent or dependent value distributions and parameter values such that late bidding may occur as a best response on Amazon. What we show, however, is how the automatic extension rule makes late bidding more difficult to achieve at equilibrium, and in particular how it rules out the sniping equilibria in those situations, identified in Section II, in which late bidding is particularly beneficial on eBay.¹⁵

III.1 A strategic model of the Amazon bidding environment

The rules of the auctions on Amazon are like those on eBay except for having a ‘soft close’ (as explained in the introduction) rather than a ‘hard close’. So, our model of Amazon will be like our eBay model except for the ending rule and, therefore, the times at which bids can be

¹⁴ The fact that different closing rules create different opportunities for timely bidding has also been observed in the design of the simultaneous ascending first-price auctions for the sale of radio spectrum licenses in the United States. See Milgrom (2003) for some related discussion of the activity rules employed to prevent sniping there.

submitted, since Amazon auctions can potentially have arbitrarily many extensions. The strategic significance of Amazon’s automatic extension rule is that there is no time at which a bidder can submit a bid to which others will not have an opportunity to respond.

- The times t at which a bid can potentially be made are:¹⁶

$$[0,1) \cup \{1\} \cup (1,2) \cup \{2\} \cup \dots \cup (n-1,n) \cup \{n\} \cup \dots$$

As in the eBay model presented earlier, a bidder may submit only finitely many bids (see footnote 5), and may *react* to a bid at time t only strictly later than t , i.e., the information sets in each interval are as in the eBay model.

- The auction extends past any time $\{n\}$ only if at least one bid is successfully transmitted at time $t = n$, and ends at the first n at which no bid is successfully transmitted.
- Bids submitted before time $t = 1$, and in any open interval $(n - 1, n)$, are successfully transmitted with certainty.
- Bids submitted at any $t = n$ for $n = 1, 2, \dots$, are successfully transmitted with $p < 1$.¹⁷

III.2 Auctions with private values on Amazon

First we emphasize that in the Amazon auction also, bidders have no dominant strategies. This is a simple corollary of the proof of the same result for the eBay model; the only change needed in the proof is to specify that the strategies considered for the other bidder have him make no bids at $t > 1$. Unlike in eBay auctions, however, late bidding (that is, bidding at $t = 1$) cannot help against incremental bidding (as defined in Section II.2) on Amazon. The reason is that

¹⁵ Because there are no dominant strategies in either our eBay or Amazon models, more general conclusions would have to specify how behavior depends on the distribution from which values are drawn and the auction parameters, and this appears potentially quite difficult.

¹⁶ We model the auction extensions as beginning with open intervals, in order to more cleanly reflect the way Amazon extensions work. However, since we will prove a theorem of the form “no equilibrium exists such that...” it is worth noting that the proof would go through even if we began an auction interval with a semi-closed interval of the form $[n - 1, n)$. That is, the lack of equilibria at which auctions are extended is not an artifact of the fact that there is no earliest possible bid in an extension of the auction.

¹⁷ A more realistic model would have auctions extended incrementally further whenever a bid was placed during an extension period, and it would not take the times $t = 1, 2, \dots, n, \dots$ be both the times in which extensions will be initiated and the times at which bids cannot be transmitted with certainty. We introduce the simpler assumptions here for clarity and simplicity; our results do not depend on these assumptions.

incremental bidders in Amazon auctions, who are prepared to continually raise their bids to maintain their status as high bidder, can be provoked to respond whenever a bid is successfully placed. In particular, incremental bidders also have time to respond to a bid that is successfully submitted at $t = 1$, so the gains from sniping against incremental bidding vanish on Amazon. But since bids at $t = 1$ run the risk of being lost (with probability $1 - p$) the costs of sniping persist.

Moreover, the following theorem indicates that, unlike in eBay style auctions with a hard close, implicit collusion by means of ‘stochastic selection’ of bids at $t = 1$ cannot easily be achieved among rational bidders on Amazon. For the theorem we assume, analogously to the proof of the corresponding eBay theorem, that all bidders have the same value H (with $H > m + 2s$, drawn from a known degenerate distribution with all its mass at H) because in this case early bidding wars yield zero profits for all bidders, winner and losers, so that the incentive for collusive behavior is particularly strong. We also assume here a “willingness to bid” in the case of indifference: A bidder who earns 0 prefers to do it by bidding v_j , and winning the auction and earning $v_j - v_j = 0$, rather than by losing the auction. (Note that while this assumption was not needed for the other results in this paper, it would not have altered them.)¹⁸

Theorem [No sniping on Amazon: an example with identical private values]. In any perfect Bayesian equilibrium in undominated strategies of our Amazon auction with identical private values, the auction is not extended (no bid comes in at $t = 1$).

Sketch of proof. At a perfect Bayesian equilibrium (in fact, subgame perfect, since the distribution of values is degenerate) in undominated strategies:

1. No bidder ever bids above his value: Any strategy that calls for bidder j to bid above H at any time t is dominated by the otherwise identical strategy in which j bids at most H at time t .
2. There is a number n^* such that the auction must end by stage $(n^* - 1, n^*) \cup \{n^*\}$, since submitted reservation prices must rise by at least s with each new submission, and (from the

¹⁸ This is a very 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.

previous paragraph) no bidder will ever submit a reservation price greater than H . If the auction gets to this last possible stage, there is only room for the price to rise by no more than $2s$.

3. If the auction gets to stage $(n^* - 1, n^*) \cup \{n^*\}$ with the current price no higher than $H - s$, the price will rise to H at some $t < n^*$, i.e. at a time when $p = 1$. To see why, assume first that the price when entering the final stage is smaller than $H - s$. Since each bidder wants to submit a (final) bid that is just marginally above the highest bid of the competitors but does not exceed H (see Section II.2), the incentive structure in this case resembles a simple Bertrand game. So, analogously to the Bertrand market analysis, in equilibrium of the subgame starting at $t = n^* - 1$ at least two bidders will bid H . As a consequence equilibrium payoffs will be zero, regardless of whether one wins or loses the auction. Now suppose the price when entering the final stage equals $H - s$. Then the willingness-to-bid assumption comes into play that states that bidders prefer to receive zero payoffs by winning the auction than by losing the auction. In equilibrium, a bidder who is not the current high bidder entering the final stage can only win if he bids H in this stage, so that in this case all bidders bid their values yielding zero profits for all bidders. Since any strategy profile that caused a player to bid at $t = n^*$ would have a lower expected payoff, because $p < 1$, than a strategy at which he bid earlier, when $p = 1$, all bids in the final stage are submitted at $t < n^*$. Thus, if the auction reaches the stage $(n^* - 1, n^*) \cup \{n^*\}$ all bidders receive zero payoffs with probability 1.
4. Inductive step. Suppose at some stage $(n - 1, n) \cup \{n\}$, it is known that at the *next* stage a price war yields zero payoffs for all bidders. Then all bidders will bid before $t = n$. (Since a price war will result if the auction is extended by a successful bid at $t = n$, any strategy profile that calls for a bidder to bid at $t = n$ is not part of an equilibrium, since that bidder gets the same or a higher expected return by bidding at $t < n$, and, by the willing bidder assumption, prefers to bid.)
5. The auction ends in the first stage: all bidders bid their true value before $t = 1$.

Strictly speaking we can conclude all bids are at $t = 0$, but taking this conclusion too literally runs the risk of over interpreting the model. Recall that we assume that one of the bids made at $t = 0$ will be randomly selected to be the first bid received; a sensible way to interpret this is that at the equilibrium bidders bid when they first notice the auction, and this is at a random time after the auction opens.

III.3 Auctions with dependent values on Amazon

As already noted, there are many ways in which bidders can have dependent values, in the sense that one bidder's information conveys information to other bidders about the value to them of the object being auctioned. Some of these could well lead to equilibria in which there is late bidding in Amazon style auctions. However, as before in the case with private values, we will show here that some kinds of late bidding that occur at equilibrium on eBay will not occur at Amazon equilibria. For this purpose, we concentrate on the "expert/dealer" model described earlier.

Theorem [An example of auctions with dependent values on Amazon]. In the "expert/dealer model" (for the one-time auction of a single item), at a perfect Bayesian equilibrium in undominated strategies, the dealer never wins an auction that uses the automatic extension rule.

Sketch of proof. The dealer values Fakes at 0, and values Genuines less than the uninformed bidder. If the object is Fake, no dealer buys it, since a strategy at which the dealer bids a positive amount for a Fake is dominated by an otherwise identical strategy at which he does not. So if a dealer bids at some time t , he reveals the object is Genuine, and at equilibrium is subsequently outbid with $p = 1$ by the uninformed bidder.

The theorem shows the sense in which Amazon's automatic extension rule robs experts of the ability to fully utilize their expertise without having it exploited by other bidders, by preventing them from sniping high quality objects at the last moment. The next section will show that significantly more late bidding is found in antiques auctions than in computer auctions on eBay, but not on Amazon. This suggests that bidding behavior responds to the strategic incentives

created by the possession of information, in a way that interacts with the rules of the auction as suggested by our models.

IV. Late bidding on eBay and Amazon

IV.1 Data description

In this section we analyze bidding history data of completed Amazon and eBay auctions.¹⁹ We downloaded data from both auction houses in two categories, “Computers” and “Antiques.” In the category Computers, information about the retail price of most items is in general easily available, in particular, because most items are new. The difference between the retail price and each bidder’s willingness to pay, however, is private information. In the category Antiques, retail prices are usually not available and the value of an item is often ambiguous and sometimes require experts to appraise. As a consequence, the bids of others are likely to carry information about the item’s value, allowing the possibility that experts may wish to conceal their information, as discussed earlier. As we will see, our data support this conjecture. In total, the data set consists of 480 auctions, 120 for each auction house-category combination, with in total 2279 bidders (740 bidders in eBay-Computers, 595 bidders in Amazon-Computers, 604 bidders in eBay-Antiques, and 340 bidders in Amazon-Antiques).

For each eBay-bidder we recorded how many seconds before the fixed deadline the last bid was submitted. (If the bid came in before the last 12 hours of the auction, we just count this bid as ‘early’). Due to Amazon’s soft close, bids within the last 10 minutes with respect to the actual deadline are not possible on Amazon. So, in order to make the timing of bids across auction houses comparable, we computed for each last bid on Amazon the number of seconds before a

¹⁹ The data are publicly provided by the auction houses. A description of some aspects of the data analyzed in this section appeared in Roth and Ockenfels (2002), but the analyses presented here are new, as is the dataset analyzed in Section V. The data set consists of randomly selected auctions completed between October 1999 and January 2000. Auctions were only included if they attracted at least two bidders, and auctions with a hidden reserve price were only considered if the reserve price was met. “Dutch Auctions”, in which bidders compete for multiple quantities of a single item under different bidding rules than described above, and “Private Auctions”, in which the identities and feedback numbers of individual bidders are not revealed, were also excluded. Amazon offers two more auction options: “10% off 1st bidder”-auctions, in which the first bidder locks in a 10% discount from the seller in case of winning, and “Take-it-price”-auctions, in which a bid equal to a seller-defined take-it-price immediately halts the auction and is accepted. (eBay’s “Buy-it-now”-option was only introduced after our data were collected.) 10%-off-1st-bidder-auctions change the incentives for the timing of bidding and were therefore excluded. Take-it-price auctions were only included if the Take-it-price was not met so that the auction ended at the initially posted deadline or, if extended, later.

‘hypothetical’ deadline. This hypothetical deadline is defined as the current actual deadline at the time of bidding under the assumption that the bid in hand and all subsequent bids were not submitted. Suppose, for example, one bid comes in 2 minutes before the initial closing time and another bidder bids 6 minutes later. Then, the auction is extended by 14 minutes. The first bid therefore is submitted 16 minutes and the second bid 10 minutes before the actual auction close. The bids show up in our data, however, as 120 and 240 seconds (before the hypothetical deadline), respectively.²⁰

We also collected information about each bidder’s feedback number. On eBay, buyers and sellers have the opportunity to give each other a positive feedback (+1), a neutral feedback (0), or a negative feedback (– 1). The cumulative total of positive and negative feedbacks is what we call the “feedback number” on eBay. Amazon provides a related though slightly different reputation mechanism. Buyers and sellers are allowed to post 1-5 star ratings about one another. We refer to the cumulative number of ratings as the “feedback number” on Amazon. In both auction houses the feedback numbers reflect the number of transactions a bidder participated in, and thus the corresponding bidders’ experiences with the auction platforms.²¹ The measure is not perfect, though, because feedback is not always given. Also, in principle, the sum of positive and negative feedbacks could cancel out, but in fact negative feedback is extremely rare. Resnick and Zeckhauser (2002) found in a large eBay data set that feedback is provided more than half the time, and less than one percent of all feedbacks given on eBay are negative. (In our eBay sample no bidder –but two sellers– have a negative feedback number.) We therefore think that it is safe to conclude that feedback numbers can serve as rough approximations for experience.

²⁰ Since we only observe the timing of last bids, this calculation implicitly assumes that there is no individual multiple bidding later than 10 minutes before the initial deadline. Late multiple bidding, however, would only strengthen our comparative results of late bidding on Amazon and eBay, because this would imply that the actual hypothetical deadline, when the last bid is submitted, is later than our computed hypothetical deadline, so that late multiple bidding would lead us to overestimate the extent of late bidding on Amazon.

²¹ Of course, the feedback numbers’ primary task is to measure the trustworthiness of the market participants (see, e.g., Bolton et al., 2003).

IV.2 Number of bidders per auction and feedback numbers

Different auction categories attract different numbers of bidders with different feedback numbers. Figure 1 shows the cumulative distributions of numbers of bidders across auction houses and categories in our sample. While most auctions had only two bidders (the minimum number that we allowed in our sample), three auctions attracted more than 14 bidders. The most popular item in our data is a Laptop auctioned on eBay with 22 bidders. The figure illustrates that eBay attracts more bidders than Amazon (two-sided χ^2 -test, $p = 0.000$) and that within each auction house Computers attract more bidders than Antiques ($p = 0.090$ for eBay and $p = 0.000$ for Amazon).

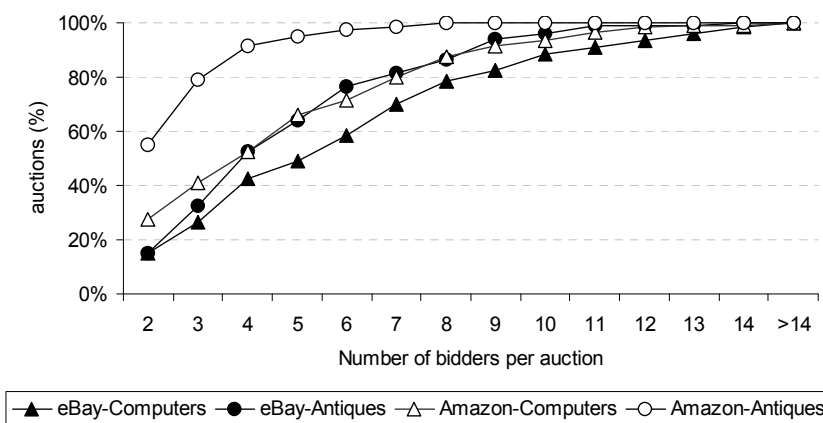


Figure 1. Cumulative distributions of numbers of bidders per auction

There is a great deal of heterogeneity with respect to bidders' feedback numbers in our data. While one third of the bidders have a feedback number of zero, the maximum number is 1162. Figure 2 shows the distributions of Amazon- and eBay-feedback numbers within categories. It illustrates that the numbers are in general higher on eBay than on Amazon (two-sided χ^2 -tests based on the categorized feedback numbers as shown in Figure 2 yield $p = 0.000$ for both, computers and antiques separately). Furthermore, antiques-bidders have higher feedback numbers than computers-bidders within each auction house ($p = 0.000$ for both, eBay and Amazon separately).




Figure 2. Distributions of feedback numbers

IV.3 Regression analyses

Table 1 shows that a considerable share of the last bids of all bidders is submitted in the very last hour of the auctions (that usually run for several days.) However, late bidding is substantially more prevalent on eBay than on Amazon. The table reveals, for instance, that on eBay 13 percent of all bidders submit their bid in the last five minutes, while few bids – only about 1 percent – come in equally late on Amazon. The difference is even more striking on the auction level: half of all eBay-Antiques auctions as compared to about 3 percent of Amazon-auctions have last bids in the last 5 minutes. The pattern repeats in the last minute and even in the last ten seconds.²²

²² The difference in late bidding was first noted in Roth and Ockenfels (2002), but space limitations did not permit any statistical or theoretical analyses of the phenomenon there. We note here that we also find scale independence in the distribution over time of bidders' last bids: The distribution of bidders' last bids within any end-interval of the auction is the same regardless of the length of the end interval, from, e.g., the last ten hours to the last ten minutes, with a high percentage of the bids concentrated at the end of the interval. This is strikingly similar to the 'deadline effect' observed in bargaining (see Roth et al., 1988, and Kennan and Wilson, 1993).

	share of <i>all</i> last bids in ...		share of auctions' last bids in ...	
	eBay	Amazon	eBay	Amazon
last hour	20 %	7 %	68 %	23 %
last 10 minutes	14 %	3 %	55 %	11 %
last 5 minutes	13 %	1 %	50%	3 %
last 1 minute	8 %	0.1 %	37 %	0.4 %
last 10 seconds	2 %	0 %	12 %	0 %

Table 1. Frequencies of late bidding on eBay and Amazon²³

The following regression analyses shed light on how late bidding relates to the numbers of bidders and feedback numbers across auction houses and categories. The dependent variable in the regression analyses in Table 2 is a binary variable with value 1 if the bidder's last bid comes in within the last ten minutes (with respect to the actual or, on Amazon, the hypothetical deadline) and 0 otherwise. The explanatory variables include binary (0/1) variables for the auction house and the item category (*eBay* and *Computers*), as well as the bidders' feedback numbers and the number of active bidders in the corresponding auction (*Feedback#* and *#bidders*). Table 2 reports the maximum likelihood probit coefficient estimates on both the bidder (all bidders) and the auction (last bidders only) level.

The regressions strongly confirm the prediction of the theoretical models that late bidding is more common on eBay than on Amazon; the coefficients for *eBay* are highly significant.²⁴ eBay auctions attract not only later bids on the bidder level (Run 1), but also the last bid of an eBay-auction has a higher probability of being late (Run 3).

²³ As explained above, the timing of bids on Amazon is defined with respect to a 'hypothetical' deadline. Also, recall that the last submitted reservation price does not necessarily win. In our eBay data, for instance, 132 final bids but only 106 winning bids were submitted within the last ten minutes.

²⁴ The results reported here pass various robustness tests. We ran Probit as well as Logit and OLS regressions using 5, 10, and 15 minutes thresholds for late bidding, and also using the classifications of feedback numbers as shown in Figure 2. We found that our qualitative findings are quite insensitive to the details of the statistical model. Also, in a pilot study, we downloaded data from eBay and Amazon in 320 auctions of computer monitors and antique books. The data set is less complete since only last bidders and only two feedback categories were considered. To the extent we can compare the data with the data reported in this paper, however, they agree in essentially all qualitative features described here.

Explanatory variables	All bidders' last bids		Auctions' last bids	
	<i>Run 1</i>	<i>Run 2</i>	<i>Run 3</i>	<i>Run 4</i>
	coeff. (<i>p</i> -value)	coeff. (<i>p</i> -value)	coeff. (<i>p</i> -value)	coeff. (<i>p</i> -value)
<i>Constant</i>	- 1.567 (0.000)	- 1.454 (0.000)	- 1.404 (0.000)	- 1.240 (0.000)
<i>eBay</i> (= 1 if eBay and 0 if Amazon)	0.773 (0.000)		1.140 (0.000)	
<i>Computers</i> (= 1 if Computers and 0 if Antiques)	- 0.187 (0.021)		- 0.260 (0.063)	
<i>Feedback#</i>	0.001 (0.012)		0.003 (0.017)	
<i>#bidders</i>	- 0.036 (0.003)		0.074 (0.001)	
<i>eBay*Computers</i>		0.477 (0.026)		0.682 (0.026)
$(1 - eBay)*Computers$		- 0.264 (0.168)		- 0.268 (0.284)
<i>eBay*(1 - Computers)</i>		0.689 (0.001)		1.033 (0.001)
<i>eBay*Feedback#</i>		0.001 (0.008)		0.003 (0.015)
$(1 - eBay)*Feedback\#$		- 0.034 (0.059)		- 0.035 (0.095)
<i>eBay*#bidders</i>		- 0.039 (0.003)		0.071 (0.007)
$(1 - eBay)*\#bidders$		- 0.026 (0.401)		0.068 (0.124)
Log-likelihood	- 650.6	- 646.4	- 239.9	- 236.8
Number of observations	2279	2279	480	480

The dependent variable is 1 for bidders whose last bid came in the last 10 minutes before the (on Amazon: 'hypothetical') deadline and 0 otherwise. The table reports maximum likelihood probit coefficient estimates.

Table 2. Probit-estimates for late bidding on eBay and Amazon

The fuller specifications in Runs 2 and 4 look at interaction effects between auction houses and categories.²⁵ Within the auction houses, the coefficient estimates in Runs 2 and 4 for *eBay*(1 - Computers)* are larger than for *eBay*Computers* indicating that eBay-Antiques trigger

²⁵ Runs 2 and 4 both include a constant variable so that we did not include a variable $(1 - eBay)*(1 - Computers)$ in order to avoid linear dependencies among the explanatory variables.

more late bidding than eBay-Computers. The difference is statistically significant ($p < 0.05$ for Runs 2 and 4 separately). A corresponding test for Amazon auctions yields no significant result ($p = 0.53$ on the bidder level and $p = 0.28$ on the auction level). Across auction houses, on the other hand, eBay attracts highly significantly more late bidding both in Computers and Antiques as well as on the bidder and the auction level ($p < 0.01$ in all four cases).

The regressions reveal an interesting correlation between feedback numbers and late bidding. Runs 2 and 4 show that the impact of the feedback number on late bidding is significantly positive on eBay ($p = 0.008$ for Run 2 and $p = 0.015$ for Run 4) and negative on Amazon ($p = 0.059$ for Run 2 and $p = 0.095$ for Run 4). This suggests that more experienced bidders on eBay go later than less experienced bidders, while experience on Amazon has the opposite effect.²⁶

In sum, there is substantially more last minute bidding on eBay than on Amazon regardless of the category. This difference also grows with the experience of the bidders. We therefore conclude that last minute bidding is not simply due to naïve time-dependent bidding. Rather, it responds to the strategic structure of the auction rules in a predictable way. In addition, in line with the theoretical considerations, significantly more late bidding is found in antiques auctions than in computer auctions on eBay, but not on Amazon.²⁷

V. Some observations on how incremental bidding affects sniping

An analysis of the dynamics of individual bidding is not possible with the data set described in Section IV, because at the time the data were collected, Amazon and eBay only provided information about the timing of each bidder's *last* bid. This changed in October 2000 when eBay started to display *all* successfully submitted bids in the bid histories. So we collected the bid history data of 308 randomly selected (computers and antiques) auctions that were completed

²⁶ The difference between the coefficient estimates is significant on the bidder level ($p < 0.05$) and weakly significant on the auction level ($p = 0.069$). Wilcox (2000) examines a sample of eBay auctions and also finds that more experienced bidders bid later.

²⁷ The regressions also reveal an impact of the number of bidders on late bidding. Runs 1 and 2 indicate that the number of bidders is negatively correlated with the average tendency of bidding late and Runs 3 and 4 indicate that the number of bidders is positively correlated with the probability that the last bid of an auction arrives late. While Runs 2 and 4 show that this size effect is significant on eBay but not on Amazon, the differences between the eBay- and Amazon-coefficients are not significant ($p = 0.70$ on the bidder level and $p = 0.96$ on the auction level). It is clear, however, that the timing cannot be explained by the 'naïve' hypothesis that more bidders per auction cause last bids to be later. In fact, last bids in eBay-Computers come earlier than last bids in eBay-Antiques, while the number of bidders per auction is actually significantly higher in eBay-Computers.

in October 2000 on eBay, with in total 1339 bidders submitting 2535 bids. Figure 3 shows the distribution of the number of bids per bidder. While most bidders (62 percent) bid only once, many bidders increased their bid in the course of the auction. The average number of bids per bidder is 1.89.

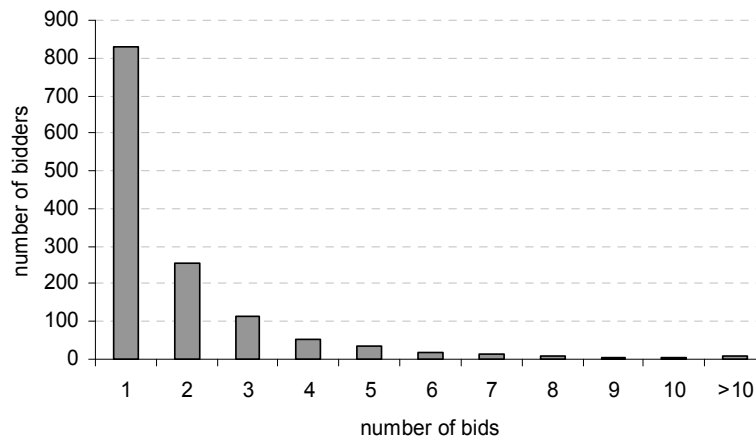


Figure 3. Distribution of number of bids submitted by a bidder

The OLS estimates in Table 3 shed some light on what triggers incremental bidding.

Explanatory variables	coefficients
<i>Constant</i>	2.043 ($p = 0.000$)
<i>Antiques</i> (= 1 if Antiques and 0 if Computers)	- 0.059 ($p = 0.550$)
<i>feedback #</i>	- 0.001 ($p = 0.000$)
<i># bidders</i>	- 0.139 ($p = 0.000$)
number of other bidders in the same auction who bid at least twice	0.349 ($p = 0.000$)
R-squared	.071
Number of observations	1339

The dependent variable is the number of submitted bids per bidder (if this number is greater than 1, then the corresponding bidder increased his reservation price at least once). The table reports ordinary least squares estimates.

Table 3. Determinants of incremental bidding

The regression shows that the number of individually submitted bids is significantly *negatively* correlated with the feedback number, suggesting that incremental bidders, unlike late bidders, are relatively less experienced.²⁸ Furthermore, the regression indicates that, controlling for the number of bidders,²⁹ the number of bids per bidder is increasing in the number of other active bidders who bid multiple times. This suggests that incremental bidding may induce bidding wars with like-minded incremental bidders.³⁰ Facing such bidding war behavior, sophisticated, experienced bidders then have an incentive to bid late to avoid pushing the price up (see Section II.2). Naive English-auction bidders, however, may also have an incentive to come back to the auction close to the deadline in order to check whether they are outbid. Indeed, a substantial share of both, one-bid bidders (12 percent) and multiple bidders (16 percent) bid as late as in the last ten minutes of the eBay auctions in the more recent sample (in the data sample discussed in Section IV, on average 14 percent of all bidders submitted last ten minutes-bids). However, the data indicates that among these late bidders, one-bid bidders submit their bid later than incremental bidders (Mann-Whitney U test, $p = 0.002$), as illustrated in Figure 4 that shows the cumulative distribution functions of last bids within the last ten minutes for bidders with one bid and multiple bids.³¹

²⁸ The average rank of the feedback of a multiple bidder is 599 (out of 1339 bidders), that of a one-bid bidder is 714 (Mann-Whitney U test, $p = 0.001$). The average rank of the feedback of a last ten minutes-bidder is 753, that of all other bidders is 654 ($p = 0.000$). (We report ranks as these are less sensitive to outliers than the actual feedback numbers themselves. The overall average feedback number is 80. The average feedback number of a multiple bidder is 61, that of a one-bid bidder is 92. The average feedback number of a last ten minutes-bidder is 99, that of all other bidders is 76.)

²⁹ The more bidders are active, the less multiple bidding is observed. Since the average bidder is outbid earlier in an auction with many bidders, he has less opportunity to bid multiple times, which may explain this number effect. Also, it takes only two to drive up the price, so not all incremental bidders have to contribute to ‘finding the right price’ early in the auction with many bidders.

³⁰ The regression also reveals that there is no difference with respect to multiple bidding between antiques and computers. In fact, while most of the results derived in Section IV apply also to the new eBay data, there is one exception: there is no statistically significant difference between antiques and computers with respect to late bidding in the new eBay data. One possibility is that this may reflect that eBay launched “a special area of eBay where you can find works of art and collectibles from leading auction houses and dealers around the world.” (<http://pages.gc.ebay.com/help/index.html>). This site, ebaypremier.com, may have drawn off the attention of experts/dealers from the antiques-categories on eBay.com, where we downloaded the data, and therefore diluted the category-effect.

³¹ Disaggregation of incremental bidding reveals a more detailed picture. Last bids that are placed after the previous bid is outbid by automatic proxy bidding tend to be earlier than last bids that are placed after the previous bid is outbid by a new proxy bid, probably reflecting that a bidding war against a human bidder takes more time than against eBay’s bidding agent. Second, only incremental bidding that pushes up the current price is negatively

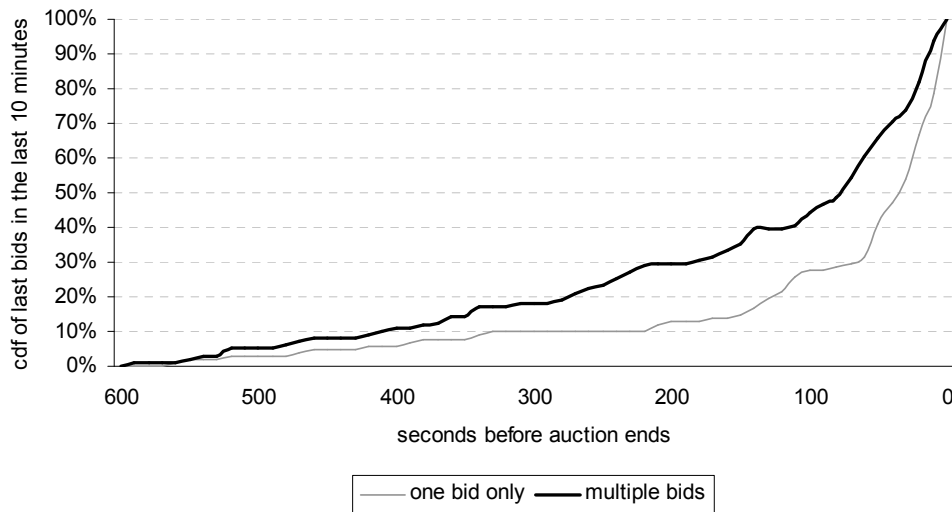


Figure 4. Cumulative distribution functions of last bids within the last ten minutes for bidders with one bid and for bidders with multiple bids

Overall, our analysis suggests that last-minute bidding is likely to arise in part as a response by sophisticated bidders to unsophisticated incremental bidding.

VII. Discussion and Conclusions

Late bidding in internet auctions has aroused a good deal of attention (more recent papers include Asker et al., 2003, Rasmusen, 2003, Schindler, 2003, and Wang, 2003). In an eBay style auction with a hard close, late bidding may be a best response to a variety of incremental bidding strategies, but sniping can also be supported as rational equilibrium behavior in both private value and common value auctions with a hard close. Thus there is no reason to believe that late bidding behavior will diminish over time, even if it initially arises as out-of-equilibrium behavior. The advantage that sniping confers in an auction with a fixed deadline is eliminated or greatly attenuated in an Amazon style auction with an automatic extension. In such an auction, an attentive incremental bidder can be provoked to respond whenever a bid is placed. So this

correlated with experience, those bidders who increase their proxy bid as the current high bidder have a similar average feedback number as one-bid bidders.

source of advantage from bidding late is reduced, while it remains costly to delay one's bid until so late that there is some probability that there will not be time to successfully submit it.

The clear difference in the amount of late bidding on eBay and Amazon is strong evidence that rational strategic considerations play a significant role. This evidence that rational considerations are at work is strengthened by the observation that the difference is even clearer among more experienced bidders and in markets in which expertise plays a role in appraising values.

Ariely et al. (2002) test the ending rules applied by eBay and Amazon in a private value environment in the laboratory. Controlling for the distribution of (induced) values, number of bidders, multiple items offered simultaneously, heterogeneity of sellers and bidders, etc., they confirm that the ending rules are sufficient to cause the differences in the timing of bids and the differential effects of experience on late and multiple bidding observed in the field.

Internet auctions and other electronic markets are an especially fertile ground for research in such design issues, because the rules of electronic markets are often explicitly given, and are easily changeable and implementable. These characteristics make electronic markets an attractive domain on which can be brought to bear a variety of complementary tools for market design, such as game theory, experimental economics, field studies, and computations (cf. Roth, 2002). Electronic markets therefore seem likely to play a major role in shaping the emerging economics of design.

References

- Ariely, Dan, Axel Ockenfels and Alvin E. Roth. "An Experimental Analysis of Ending Rules in Internet Auctions." Working paper, Harvard University, 2002 (available at: <ftp://ftp.mpiew-jena.mpg.de/pub/esi/staff/ockenfels/aor.pdf>).
- Asker, John, Brit Grosskopf, C. Nicholas McKinney, Muriel Niederle, Alvin E. Roth and Georg Weizsäcker. "Teaching Auction Strategy using Experiments Administered via the Internet." Working paper, Harvard University, 2003.
- Bajari, Patrick and Ali Hortaçsu. "Winner's Curse, Reserve Prices and Endogenous Entry: Empirical Insights from eBay Auctions." *Rand Journal of Economics*, forthcoming.
- Bolton, Gary E., Elena Katok, and Axel Ockenfels. "How Effective are Online Reputation Mechanisms? An Experimental Investigation." Working paper, Smeal College of Business Administration, Penn State University, 2003.
- Kennan, John and Robert Wilson. "Bargaining with Private Information." *Journal of Economic Literature*, 1993, 31(1), 45-104.
- Landsburg, Steven E. "My Way to the eBay: An Economics Professor's Strategy for Winning Online Auctions," *Slate*, Thursday, April 8, 1999 (available at: <http://slate.msn.com/Economics/99-04-08/Economics.asp>).
- 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, 2002.
- Milgrom, Paul. *Auction Theory for Privatization*. In preparation, 2002.
- 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.
- Rasmusen, Eric. "Strategic Implications of Uncertainty Over One's Own Private Value in Auctions." Working paper, Indiana University, 2003 (available at: <http://php.indiana.edu/~erasmuse/papers/auction.pdf>).
- 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." *Econometrica*, 70, 4, July 2002, 1341-1378.
- 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.
- Schindler, Julia. "Late Bidding on the Internet." Working paper, University of Vienna, 2003.
- Vickrey, William. "Counterspeculation, Auctions, and Competitive Sealed Tenders." *Journal of Finance* 1961, 16, 8 -37.
- Wang, Tao-yi Joseph. "Is Last Minute Bidding Bad?" Working paper, UCLA, 2003.
- Wilcox, Ronald T. "Experts and Amateurs: The Role of Experience in Internet Auctions." *Marketing Letters*, 2000, 11(4), 363-374.