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

Axel Ockenfels & Alvin E. Roth*

December 20, 2001

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 private value auctions, as well as in common value auctions Sniping in fixed-deadline auctions also arises away from 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. Finally, we find scale independence in the distribution over time of bidders' last bids, of a form strikingly similar to the 'deadline effect' noted in bargaining. (*JEL* C73, C90, D44)

---

* Ockenfels: Faculty of Economics and Management, University of Magdeburg, P.O. Box 4120, D-39016 Magdeburg, Germany; e-mail: axel.ockenfels@ww.uni-magdeburg.de (URL: http://www.uni-magdeburg.de/vwl3/axel.html). Roth: Harvard University, Department of Economics and Graduate School of Business Administration, Soldiers Field Road, Baker Library 183, Boston, MA 02163, USA; e-mail: aroth@hbs.edu (URL: http://www.economics.harvard.edu/~aroth/alroth.html).
We gratefully acknowledge helpful conversations on this subject with Estelle Cantillon, Scott Cook, Jeff Ely, Ed Glaeser, Werner Güth, Seungjin Han, Ehud Kalai, Bertrand Koebel, Eric Maskin, Muriel Niederle, Martin Osborne, Ariel Pakes, Jack Porter, Jean-Francois Richard, Uri Rothblum, Hal Varian, and comments from audiences at Augsburg, Berkeley, Berlin, Bielefeld, Bilbao, Bonn, Columbia, Dortmund, Erfurt, Frankfurt, Gerzensee, Harvard, Koblenz, LSE, Magdeburg, Minnesota, Munich, Northwestern, Stanford, and Wellesly. We also thank Debora Y. Lewi and John Wooders for correcting mistakes in an earlier version, the many bidders who allowed us to interview them, and the readers of a New York Times column by Hal Varian and a Wall Street Journal article by Joel Rosenblatt that mentioned an earlier version of this paper for many stimulating opinions. This project received financial support from the NSF, the Harvard Business School, and the DFG 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. The auction interface typically allows a bidder to submit a reservation price, and to have this (maximum) price used to bid for him by proxy. That is, a bidder can submit a reservation price (called a proxy bid) early in the auction and have the resulting bid register as the minimum increment above the previous high bid. As subsequent reservation prices are submitted, the bid rises by the minimum increment until the second highest submitted reservation price is exceeded. Hence, an early bid with a reservation price 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 reservation price.[1] 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.

In a recent paper, Roth and Ockenfels (2001) noted that there is more late bidding on eBay than 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,[2] that arise despite advice from auctioneers that bidders

---

[1] Second price auctions have been used long before the first internet auctions. Goethe designed a second price auction in 1797 to elicit his publisher's willingness to pay for his manuscript "Hermann und Dorothea" (see Moldovanu and Tietzel, 1998), and many stamp auctions used second price rules as early as in the nineteenth century (see Lucking-Reiley, 2000). In his seminal work, Vickrey (1961) offered the first theoretical analysis of the sealed-bid second price auction and compared it to other auction formats. Since then, economists often refer to the sealed-bid second price auction as a "Vickrey-auction".

[2] 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."

should simply submit their maximum willingness to pay, once, early in the auction.[3] 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, it is a *dominant* strategy for bidders to submit as their reservation price their maximum willingness to pay. In 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).[4] Malhotra and Murnighan (2000) also conclude they are seeing irrational late bidding, in a fascinating 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, 2001). Human and artificial bidders do not differ in this respect. The last-minute bidding service esnipe.com, that offers to automatically place a predetermined bid a few seconds before the end of the eBay auction, admits that it cannot make sure that all bids are actually placed: "We certainly wish we could, but there are too many factors beyond our control to guarantee that bids always get placed. While we have a very good track record of placed bids, network traffic and eBay response time can sometimes prevent a bid from being completed successfully. This is the nature of sniping" (http://www.esnipe.com/faq.asp, 2000).[5]

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, and even multiple bidding may be consistent with equilibrium, perfectly rational behavior in both public value and private value auctions. However, we will

[3] 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).

[4] There may not exist equilibria in dominant strategies, however, if the price is required to be (at most) one minimum increment above the second highest bid but cannot exceed the highest bid as it is the case on eBay (see Section II.1).

[5] Another online tool for snipers is provided by ascentive.com, that offers to synchronize the computer system clock with an atomic clock server over the internet, if, for instance, one is "Tired of losing those eBay auctions by mere seconds" (http://www.ascentive.com/products/ActiveTime-welcome.html).

also show that rational late bidding in both common and private value auctions 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 significantly 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.

The data also reveal a striking scale invariance property of the timing of last bids on both eBay and Amazon. 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.

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.[6] We will construct a model with a multiplicity of equilibria, including equilibria with last-minute bidding, even in purely private value auctions. For clarity, we first present the strategic structure of the auction, to which we will add the players' valuations (private value or common value) when we examine particular equilibria.

---

[6] 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 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.

The strategic model (the eBay game-form):

- There are n bidders, $N = \{1,\ldots,n\}$.
- There is a minimum initial bid m, and a smallest increment s by which subsequent bids must be raised.[7]
- 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 are two exceptions 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. (If 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 the current high bidder submits a new, higher reservation price, the current price is not raised, although the number of bids is incremented.[8]
- Each submitted reservation price must exceed both the current high bid and the bidder's last submitted reservation price (i.e., a bidder cannot lower his own previous reservation price).
- The bid history lists the current price and the submission time of each bidder's last submitted reservation price[9] 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

---

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.

[7] In eBay the smallest increment increases as the price increases; see http://pages.ebay.com/help/basics/e_bid_conf3.html for a table that outlines the different increment levels.

[8] The one exception is if the current price is less than the minimum increment above the second highest submitted willingness to pay, which can occur 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.

[9] 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

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}.[10] 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.

- 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.
- If multiple bids are submitted simultaneously at the same instant t, then they are randomly ordered, and each has equal probability of being received first.
- Bids submitted before time $t = 1$ are successfully transmitted with certainty.
- At time $t = 1$, the probability that a bid is successfully transmitted is $p < 1$.[11] The probability p is assumed to be exogenous, i.e. not influenced by the bidding in a given auction.[12]

II.2 Private value auctions on eBay

In a private value auction, each bidder j has a true willingness to pay, v_j, distributed according to some known, bounded distribution F. 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. 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 no dominant strategies in the eBay environment.

---

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.

[10] 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.'

[11] 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.

[12] Of course, 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.

Theorem [No dominant strategies]. A bidder in the second price private value eBay-model auction does not have any dominant strategies.

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.

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 $t < 1$ (which would provoke a counterbid of B, and cause j to win at most zero), 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.[13]

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 $t < 1$, at which the bid will be transmitted with certainty instead of with probability $p < 1$). 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, i.e., does not have any strategy that is a best reply to all strategies of the other player.

One might wonder if in the subgame starting 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 no bids were submitted at $t < 1$, and that bidder i plans to bid some $b_i < v_j - s$ at $t = 1$. In this situation bidder j would have to pay a price of $b_i + s$ if he bids his value at $t = 1$ 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. In particular, the closer bidder *j*'s bid to b_i the higher his expected payoff as long as bidder *j*'s bid exceeds b_i (recall that in case of a tie the winner is chosen randomly). So, bidder *j* wants to condition his bid on bidder *i*'s bid, implying that no dominant strategy exists in the subgame starting at $t = 1$.[14]

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. So, 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 representatively define incremental bidding as a 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, whatever comes first.[15] 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.

---

[13] Other bids at $t = 1$ would yield the same payoff against *i*'s strategy, and would also be best replies.

[14] 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, depending on the underlying value distribution, bidding values may be a Bayesian Nash-equilibrium in the subgame starting at $t = 1$ (though not in dominant strategies), and this is in fact the case in the examples that we will elaborate in this paper. Furthermore, 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" in the subgame 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.

[15] 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, 2000).

Sketch of proof. 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 secures that the bid is accepted 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 raise the best price that is achievable at $t = 1$ above $m + s$.[16]

Even though much late bidding appears to be a strategic reaction to (non-equilibrium) incremental bidding,[17] it is still of interest to know if late and multiple bidding can themselves constitute equilibrium behavior despite the risk involved in late bidding. The next theorem shows that they can, which suggests that, even if bidders initially learn to bid late in response to unsophisticated behavior of others, late bidding is likely to persist even as all bidders gain experience. The intuition behind last-minute bidding at equilibrium in a private value auction 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

---

[16] 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 auction environment such as the distributions of values and the values of auction parameters p, m and s.

[17] See Section V for field evidence and Ariely et al. (2001) for experimental evidence on incremental bidding. Anecdotal evidence can be found in Ockenfels and Roth (2001) and comes from auctionwatch.com, a rich source of information for bidders and sellers, that proposes an incremental bidding strategy: "Another bidding tactic is sentry bidding–placing a bid and then keeping close watch on it for the duration of the auction. If others try to outbid you, you can quickly place a counter-bid to regain your high bidder status. (…) This could lead to an all-out bidding war, so be clear on how much you're willing to spend for an item and how closely you can monitor the auction. And beware of the last-second snipe from your worthy opponent." (http://www.auctionwatch.com/awdaily/tipsandtactics/buy-bid.html, 2000).

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.

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. But the positive probability that a last-minute bid will fail to register also creates a cost, so there is also some incentive to bid early. Thus there can also exist equilibria at which buyers may bid both early *and* at the very last moment.

Theorem [Private value auctions on eBay]. There can exist 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. There can also exist equilibria in which bidders want to bid both early and late.

Proof. Consider the case of two bidders, $N = \{1,2\}$. For simplicity we assume that both bidders have the same value H (with $H > m + 2s$). We will first demonstrate that mutual sniping without multiple bidding may constitute an equilibrium; then we add multiple bidding. In particular, in this simple setting both bidders will earn zero at any equilibrium in which they both bid their values with certainty, but there are equilibria that give 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. 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 $t = 1$, mutually bidding true values is a Nash equilibrium. This proves that the indicated strategy is a best reply for each bidder.

Now consider the following strategies that call for each bidder to bid multiple times and that also constitute an equilibrium for all $p > 0$. On the equilibrium path, each bidder i's strategy is to bid L with $m + s \leq L < H - s$ at $t = 0$, and then to bid H at $t = 1$. Off the equilibrium path, if the price becomes greater than L at a $t' < 1$, each bidder's strategy is to bid H at some $t > t'$ such that $t < 1$. That is, each bidder's strategy is to bid L early, at $t = 0$, and do nothing else until $t = 1$, unless the price exceeds L prior to $t = 1$ in which case the strategy calls for the bidder to promptly bid H. Thus, on the equilibrium path, the price is L at all t with $t < 1$, since at $t = 0$ both players bid L, so the first and second highest bids are equal. Since both bids at $t = 0$ are submitted simultaneously, one of the bidders is selected to be the high bidder early in the auction with probability 1/2. The earnings of this bidder in equilibrium are $H - L$ if both bids at $t = 1$ are lost (which happens with probability $(1 - p)(1 - p)$), $H - L - s$ if only the other bidder's bid at $t = 1$ is lost (with probability $1 - p$), and zero otherwise. The earnings of the low bidder at $t = 0$ are $H - L - s$ if his bid at $t = 1$ is successful but the opponent's bid is lost (with probability $p(1 - p)$), and zero otherwise. No unilateral deviation from equilibrium strategies is profitable: A bid higher than L at any $t < 1$ is detected for sure, because it causes the price to rise above L, and earns zero. A bid below L at $t = 0$ eliminates the potential gain that comes from being randomly selected to be the high bidder early in the auction, and winning the auction if both bids at $t = 1$ are lost. Finally, as before, in the subgame starting at $t = 1$, bidding true values constitutes a Nash equilibrium. This proves that the indicated strategy is a best reply for each bidder.

The proof shows that even at equilibrium in private value auctions, 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. Furthermore, in an equilibrium with multiple bidding, one bidder will win the auction, even if all bids at $t = 1$ are lost. Raising one's bid in the course of an auction, as frequently observed in eBay, is therefore not inevitably a sign of irrationality or inexperience on the part of bidders, nor an indication of common value characteristics of the item being auctioned. Multiple bidding can occur at equilibrium of a private value internet auction with a hard close (though the evidence in

Section V suggests that the observed multiple bidding is at least in part due to unsophisticated, inexperienced behavior). For the multiple bidding case, the proof assumes that the high bidder at $t = 0$ is selected randomly, because both identical bids are placed simultaneously. A more realistic model would not assume that bidders are immediately aware of all new auctions on eBay. In this case, an analogous, slightly modified equilibrium strategy would call for the bidder who discovers the auction first, at some $t < 1$, to promptly bid L, and for another bidder, who notices the auction only later, at some t' with $t < t' < 1$, to refrain from bidding more than L, in order to avoid an early bidding war yielding zero profit.

A few remarks on the proof are in order. Different values of L describe different equilibria, and bidders may want to keep the early price L as small as possible so that the potential gains from sniping, $H - L$ ($- s$), are maximal. But the smallest possible bid, $L = m$, is by assumption excluded in the proof. The reason is that if both players bid m at $t = 0$, the bid that is randomly chosen to come in second would not be processed since it does not exceed the current price.[18] However, an early price of m is nevertheless feasible in equilibrium, but unlike in the case discussed above, cheating by bidding more than $L = m$ is only detected with probability 1/2 (if the bid is chosen to come in second), so that a deviation from equilibrium yields higher expected profits than in the case described in the proof above.

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 different values, the existence of equilibria with late and multiple bidding 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.

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 early 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 $1 - p^2$).

---

[18] If $L \geq m + s$, the price after the first but before the second bid of L has been processed is m, yielding a minimum bid of $m + s \leq L$, so that the second bid can be accepted and raise the price to L.

II.3 Common value auctions on eBay

The simplest models of common value auctions treat the bidders as more or less symmetric, except for their private signals about the common 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 (2000), 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 common 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.[19]

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 =

[19] Such a situation is described by esnipe.com as one motivation for why one would want to snipe: "Experienced collectors often find that other bidders watch to see what the experts are bidding on, and then bid on those items themselves. The expert can snipe to prevent other bidders from cashing in on their expertise." (http://esnipe.com/faq.asp). Some experienced antiques bidders in the survey by Roth and Ockenfels (2001) also 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 (as do the surveys in Roth and Ockenfels, 2001).

$1 - p_F$. 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 $m < H - c + s < H$. Then the common 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 [Common value auctions on eBay]. When the probability that the object is Fake is sufficiently high, there is a subgame perfect 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 $t < t' < 1$.

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 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 $t < t' < 1$. However so long as $p_F/p_G > (H - m)/m$, it never pays for the uninformed bidder to bid at $t = 1$ 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. When no player has bid at $t < 1$, bidding true value at $t = 1$ is part of a subgame perfect equilibrium strategy when the object is Genuine. So 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.[20]

[20] The equilibrium is independent of 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. However, the

Because there are many ways that values can be interdependent (as opposed to the "pure" or "mineral rights" model of common values), we do not seek here to develop a general model of common values. We simply re-emphasize that the demonstration in the previous section, that late bidding can occur at equilibrium in purely private value auctions, 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 private value case. 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.[21] We caution from the start that our examples in this section do not imply that there is no combination of independent or interdependent value distributions and parameter values such that late bidding may occur as a best response on Amazon. What the theorems clearly 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

---

game has more equilibria. For instance, if we allow weakly dominated strategies, there are 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.

[21] 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 (2000) for some related discussion of the activity rules employed to prevent sniping there.

those situations, identified in Section II, in which late bidding is particularly beneficial on eBay.[22]

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 no fixed deadline. The relevant Amazon rules are the following:

> "We know that bidding may get hot and heavy near the end of many auctions. Our Going, Going, Gone feature ensures that you always have an opportunity to challenge last-second bids. Here's how it works: whenever a bid is cast in the last 10 minutes of an auction, the auction is automatically extended for an additional 10 minutes from the time of the latest bid. This ensures that an auction can't close until 10 'bidless' minutes have passed. The bottom line? If you're attentive at the end of an auction, you'll always have the opportunity to vie with a new bidder." Amazon (2001)

Our model of Amazon will be like our eBay model except for the ending rule and, therefore, the times at which bids can be 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:[23]

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

  (As in the eBay model presented earlier, a bidder 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.

---

[22] 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 of values and the auction parameters, and this appears potentially quite difficult.

[23] 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.

- 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 $n = 1,2, \ldots,$ are successfully transmitted with $p < 1$.[24]

III.2 Private value auctions 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 $t > 1$.

Corollary [No dominant strategies]. A bidder in the second price private value internet auction with automatic extension does not have any dominant strategies.

Contrary to eBay auctions, however, late bidding (that is, bidding at $t = 1$) cannot help against incremental bidding (as defined in Section II.2) on Amazon:

Theorem [Best response to incremental bidding on Amazon]. Bidding at $t = 1$ is not a best reply to incremental bidding.

Sketch of proof. In Amazon auctions incremental bidders, 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.

---

[24] 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 $t = 1,2,\ldots, n, \ldots$ 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.

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 $H > m + 2s$) 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 earning $v_j - v_j = 0$, rather than by losing the auction. (Note that this assumption would not have interfered with any other result in this paper).[25]

Theorem [Private value auctions on Amazon]. In any subgame perfect equilibrium in undominated strategies of our Amazon private value auction with identical values, the auction is not extended (no bid comes in at t = 1).

Sketch of proof. At a subgame perfect equilibrium 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 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 at $t = n^*$ 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

[25] 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.

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 at $t = n^* - 1$ 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 at $t = n^* - 1$ can only win if he bids H in the final stage, so that in this case all bidders bid their values in the last possible stage 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 $t < n$.)
5. The auction ends in the first stage: all bidders bid their true value before $t = 1$. (There is also no multiple bidding.)[26]

III.3 Common value auctions on Amazon

As already noted, there are many ways in which bidders can have common 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 private value case, we will show here that

[26] 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.

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 [Common value auctions on Amazon]. In the "expert/dealer model" (for the one-time auction of a single item), at a subgame perfect 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[27]

IV.1 Data description

In this section we analyze bidding history data of completed Amazon and eBay auctions that are publicly provided by the auction houses.[28] We downloaded data from both auction houses in two

---

[27] A description of some aspects of the data analyzed in this section appeared in Roth and Ockenfels (2001), but the analyses presented here are new, as is the dataset analyzed in section V.

[28] The data set consists of randomly selected auctions completed between October 1999 and January 2000 that met certain selection criteria. 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.

categories, "Computers" and "Antiques."[29] 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.[30] 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.[31] For each auction, we collected data about the number of bids, number of bidders, and about whether there was a reserve price. On the bidder level, we collected information about the "timing" of the last bid and about each bidder's "feedback number". Both variables will be described in detail in the following paragraphs.[32]

Both auction houses provide information about when each bidder's last bid is submitted. For each bidder we recorded how many seconds before the deadline the last bid was submitted if the bid came in within the last 12 hours of the auction. (If the bid came in before the last 12 hours of the auction, we just count this bid as 'early'). While this information is readily provided in eBay's bid histories of completed auctions, the end time of an auction in Amazon is endogenously determined since an auction continues past the initially scheduled deadline until ten minutes have passed without a bid. We therefore computed for each last bid in Amazon the

---

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. 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.

[29] The items considered here were found in eBay's categories "Computers: Hardware: Monitors" and "Notebooks", respectively, in eBay's category "Antiques: Ancient World", in Amazon's category "Computers & Software" searching for the Keywords "Monitor" and "Laptop", respectively, and in Amazon's category "Arts & Antiques" searching for the Keywords "antiques" or "ancient", respectively. After we collected the data, eBay's "Ancient World" category was renamed "Antiquities".

[30] eBay recommends using "appraisal services" such as offered by www.isa-appraisers.org if one needs help to evaluate an item.

[31] Specifically, we have 740 bidders in eBay-Computers, 595 bidders in Amazon-Computers, 604 bidders in eBay-Antiques, and 340 bidders in Amazon-Antiques.

number of seconds before a '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 one minute before the initial closing time and another bidder bids 8 minutes later. Then, the auction is extended by 17 minutes. The first bid therefore is submitted 18 minutes and the second bid 10 minutes before the actual auction close. The bids show up in our data, however, as 60 and 120 seconds (before the hypothetical deadline), respectively.[33]

In eBay, buyers and sellers have the opportunity to give each other a positive feedback (+1), a neutral feedback (0), or a negative feedback (– 1) along with a brief comment. A single person can affect a user's feedback number by only one point (even though giving multiple comments on the same participants is possible). The cumulative total of positive and negative feedbacks is what we call the "feedback number" in eBay. It is displayed in parentheses next to the bidder's or seller's eBay-identification name. Amazon provides a related though slightly different reputation mechanism. Buyers and sellers are allowed to post 1-5 star ratings about one another. Both the average number of stars and the cumulative number of ratings are displayed next to the bidder's or seller's Amazon-identification name.[34] We refer to the cumulative number of ratings as the "feedback number" in Amazon. Since in both auction houses the feedback numbers (indirectly) reflect the number of transactions, they might serve as approximations for experience and, more cautiously, as an indicator of expertise.[35]

---

[32] Actually, we collected more data such as the minimum bid and for each bidder the amount of the last submitted bid. However, since our analysis focuses on the dynamics of bidding rather than on pricing and revenue questions these data will not play a role in our analysis.

[33] 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. Suppose, for instance, that the buyer who bids last in the example above also bids 5 minutes before the *initial* closing time. Then, we would have misrepresented the timing of the other buyer's bid, since this bid was then actually submitted 6 minutes before the hypothetical deadline rather than only one minute as computed above. The potential effect of this bias is, however, very small. In total, 28 out of 240 Amazon auctions in our sample were extended. In 26 of these auctions, only one bidder and in the other 2 auctions two bidders bid within the last 10 minutes with respect to the initial deadline. Therefore, we may misrepresent the timing of up to 30 out of 935 Amazon-bidders. Note that the possible misrepresentation of timing with respect to the hypothetical instead of the actual closing time leads us, if at all, to *overestimate* the extent to which Amazon-bidders bid late. This would only strengthen our comparative results of late bidding in Amazon and eBay.

[34] By clicking on the feedback number in eBay and on the stars in Amazon, everyone has immediate access to the whole history of feedback comments.

[35] Note that the feedback number in eBay is the sum of positive and negative feedbacks. Hence, if positive and negative feedbacks were left with comparable probabilities, the feedback numbers could not be interpreted as

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 in 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).

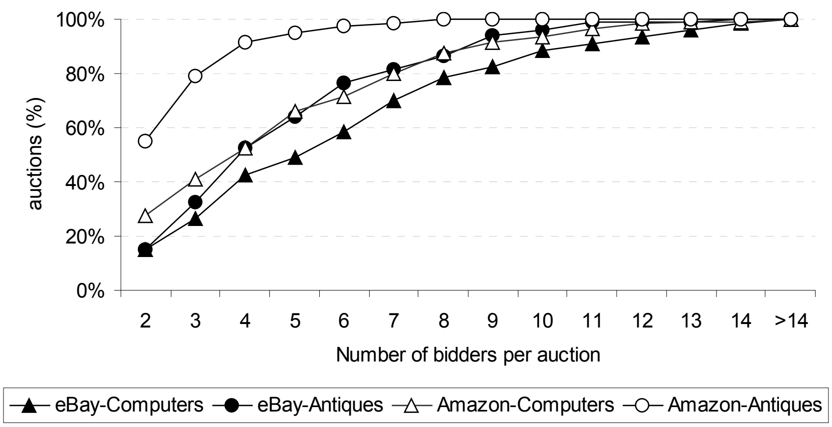

Figure 1. Cumulative distributions of numbers of bidders per auction

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

experience or expertise. The fact, however, that in our eBay sample *no* bidder (but two sellers) had a negative feedback number while about one third have zero feedback numbers indicates that negative feedbacks are left very rarely (see also Resnick and Zeckhauser, 2001, who conclude from a larger data set that feedback in eBay is "almost always" positive). In fact, eBay encourages users to avoid negative feedback numbers: "If you were treated well by a buyer or seller, reward him or her with a positive comment. If you were treated poorly, try to resolve the problem first by contacting the other person. Most problems can be corrected by improving communication between buyer 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, eBay and Amazon separately). Furthermore, antiques-bidders have higher feedback numbers than computers-bidders within each auction house ($p = 0.000$ for both, eBay and Amazon separately).

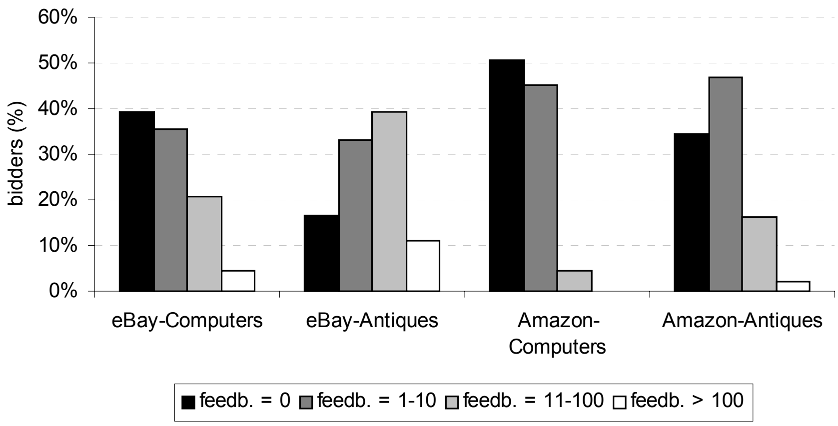

Figure 2. Distributions of feedback numbers

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 in eBay than in Amazon. The table reveals, for instance, that in eBay 13 percent of all bidders submit their bid in the last five minutes, while only few bids – about 1 percent – come in equally late in 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 tens seconds (see Section VI for an analysis of the full shape of the distributions).[36]

and seller. If things are still not resolved, you may leave a negative comment." (http://pages.ebay.com/help/basics/f-feedback.html#2, 2000).

[36] The difference in late bidding was first noted in Roth and Ockenfels (2001), but space limitations did not permit any statistical or theoretical analyses of the phenomenon there.

	share of *all* 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 in eBay and Amazon[37]

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, in 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.[38] 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).

---

[37] Recall that the timing of bids in Amazon is defined with respect to a 'hypothetical' deadline that differs from the actual closing time if a bid comes in later than ten minutes before the initial end time. Recall also that the last bidder is not necessarily the high bidder since an earlier submitted proxy bid can outbid subsequently incoming bids. In our eBay data, for instance, 132 final bids but only 106 winning bids were submitted within the last ten minutes. In Amazon the corresponding frequencies of final and winning bids are 28 and 20, respectively. We finally note here that it is not too unusual to see the auction price in eBay double in the last sixty seconds, and since it takes some seconds to make a bid, bidders attempting to submit a bid while the price is rising so rapidly may receive an error message telling them that their bid is under the (current) minimum bid. These eBay-bidders, who attempted to bid in the last minute, are not represented in these data, since their last minute bids did not register as bids in the auction.

[38] The results reported here pass various robustness tests. First, we ran Probit as well as Logit and OLS regressions using 5, 10, and 15 minutes thresholds for late bidding. We found that our qualitative findings are quite insensitive to the statistical model or the threshold for late bidding. (We discuss the full shape of the distributions of last bids in the next section.) Second, 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

The fuller specifications in Runs 2 and 4 look at interaction effects between auction houses and categories.[39] 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 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).

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.

[39] 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.

Explanatory variables	All bidders' last bids		Auctions' last bids	
	Run 1	*Run 2*	*Run 3*	*Run 4*
	coeff. (*p*-value)	coeff. (*p*-value)	coeff. (*p*-value)	coeff. (*p*-value)
Constant	– 1.567 (0.000)	– 1.454 (0.000)	– 1.404 (0.000)	– 1.240 (0.000)
eBay (= 1 if eBay and 0 if Amazon)	0.773 (0.000)		1.140 (0.000)	
Computers (= 1 if Computers and 0 if Antiques)	– 0.187 (0.021)		– 0.260 (0.063)	
Feedback#	0.001 (0.012)		0.003 (0.017)	
#bidders	– 0.036 (0.003)		0.074 (0.001)	
*eBay*Computers*		0.477 (0.026)		0.682 (0.026)
(1 – *eBay*)**Computers*		– 0.264 (0.168)		– 0.268 (0.284)
*eBay**(1 – *Computers*)		0.689 (0.001)		1.033 (0.001)
*eBay*Feedback#*		0.001 (0.008)		0.003 (0.015)
(1 – *eBay*)**Feedback#*		– 0.034 (0.059)		– 0.035 (0.095)
eBay#bidders*		– 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 (in Amazon: 'hypothetical') deadline and 0 otherwise. The table reports maximum likelihood probit coefficient estimates.

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

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 in eBay (p = 0.008 for Run 2 and p = 0.015 for Run 4) and negative in 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 in Amazon has the opposite effect, as suggested by the strategic hypotheses.[40] (Wilcox, 2000, examines a sample of eBay auctions and also finds that more experienced bidders bid later.)

In sum, there is substantially more last minute bidding in eBay than in Amazon regardless of the category. This difference also grows with the experience of the bidders. It is therefore safe to conclude that last minute bidding is not simply due to naïve time-dependent bidding.[41] 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.[42]

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 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.

---

[40] The difference between the coefficient estimates is significant on the bidder level ($p < 0.05$) and weakly significant on the auction level ($p = 0.069$).

[41] Sniping may also be a rational strategy if the act of bidding incurs an irreversible entry cost. However, if such a cost is relevant in both eBay and Amazon auctions, it could not explain the differences between eBay and Amazon that we observe.

[42] 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 in eBay but not in 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.

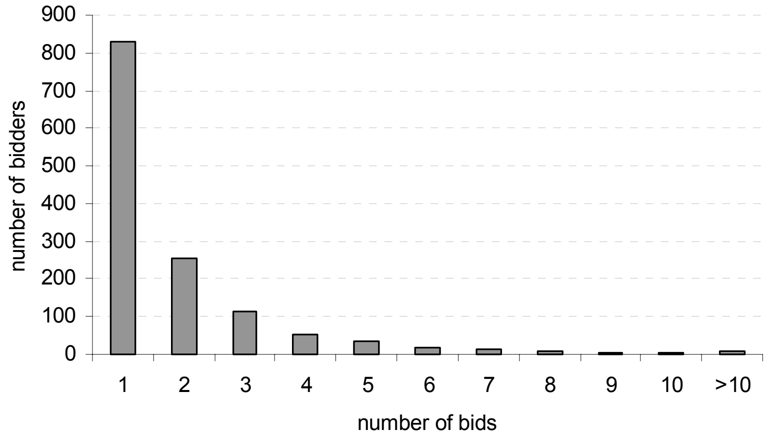

Figure 3. Distribution of number of bids submitted by a bidder

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

Explanatory variables	coefficients
Constant	2.043 (p = 0.000)
Antiques (= 1 if Antiques and 0 if Computers)	– 0.059 (p = 0.550)
feedback #	– 0.001 (p = 0.000)
# bidders	– 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.[43] Furthermore, the regression indicates that, controlling for the number of bidders,[44] 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.[45] 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 the following figure that shows the cumulative distribution functions of last bids within the last ten minutes for bidders with one bid and multiple bids.[46]

---

[43] 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 feedback number of a last ten minutes-bidder is 99, that of all other bidders is 76.)

[44] 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.

[45] 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. 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.

[46] 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 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.

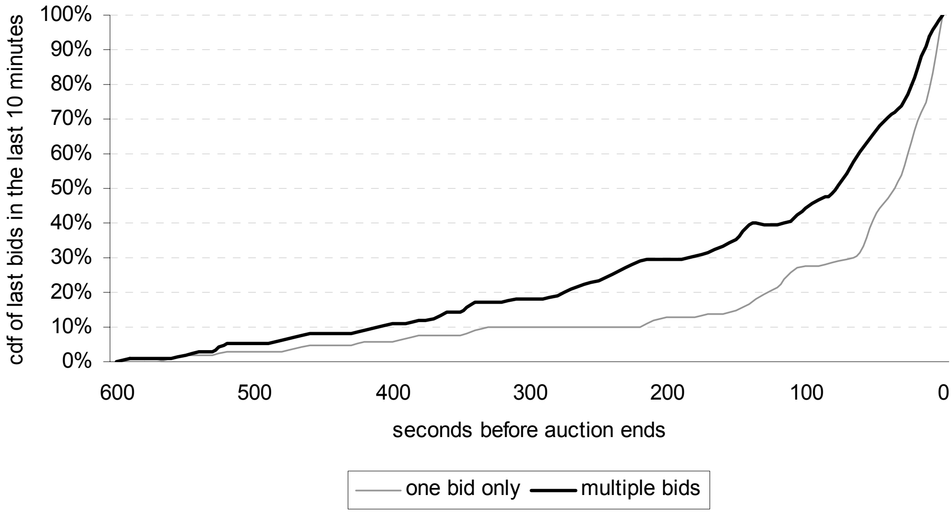

Figure 4. Cumulative distribution functions of last bids within the last ten minutes for bidders with one bid and for bidders with multiple bids

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

VI. A closer look at the timing of last bids

The simple strategic model of eBay presented in section II.1 of this paper supposes that there is a unique time t = 1, such that all bidders can respond to bids that come in before but cannot respond to bids that come in later than this threshold. This abstracts away from the fact that the ability and opportunity to bid late may differ for different bidders. Even if one presumes that all bidders actually want to bid late, the probability that they do so depends on a mixture of technical, physical, and social factors such as weekday, working hours, time zone, physical reaction time, internet congestion, modem speed, etc.

We find a striking regularity in the distribution over time of bidders' last bids: it is almost impossible to say whether a distribution of last bids is drawn from, say, the last hour or from the last 12 hours if no information about the time scale is given. Figure 5 illustrates this regularity with the help of the data set of Section IV. It shows empirical cumulative distribution functions F_T for end intervals $[0,T]$, where 0 represents the auction deadline and T = 12 hours, 6 hours, 3 hours, 1 hour, 30 minutes, 10 minutes, and, for eBay only, 1 minute. In the figures below, each

F_T graphs the percentage of last bids received in the interval [0,T] that were received in each decile of the interval. That is, the percentage of bids received in the first decile (T/10) are the percentage of the bids that were recorded in the last 10% of the interval. (Since we do not have data about the distribution of bids that are submitted more than 12 hours before the deadline, we restrict our analysis in the following to bids with $T \leq$ 12 hours.)

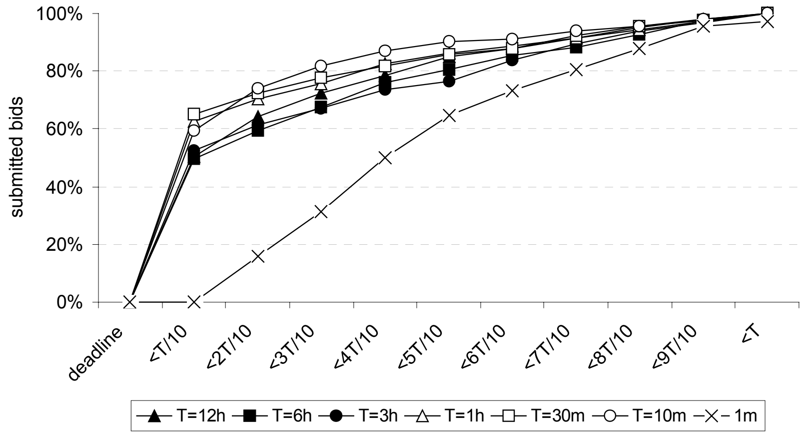

(a) eBay

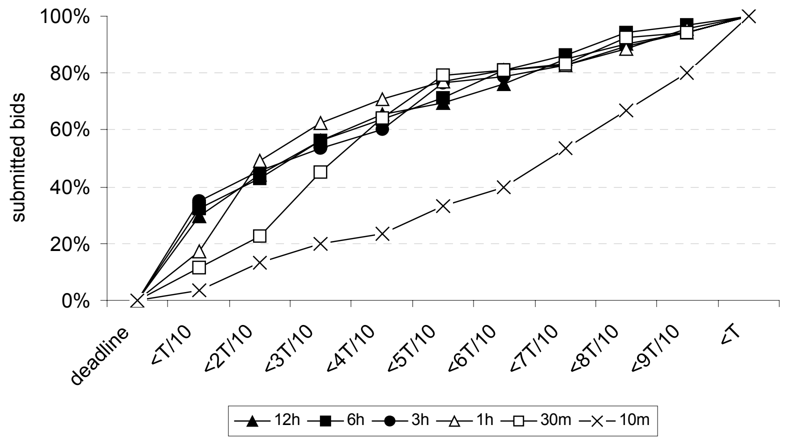

(b) Amazon

Figure 5. Scale independence in the distributions of last bids in eBay and Amazon

While shorter intervals generally induce somewhat more bidding in the first decile on eBay and less bidding in the first decile on Amazon,[47] the probability mass in each decile is surprisingly independent of the length of the corresponding interval with the exception of the shortest interval in eBay and Amazon, respectively. On eBay, averaging over all end intervals with $T > 1$ minute in Figure 5, about 57% of all bids are submitted in the first decile (with a mean deviation of 6%). That is, if you look at any end interval T, from 12 hours to 10 minutes before the end of the auction, about 57% of the last bids that arrive in that interval come in the final ten percent of the interval. The very last minute on eBay does not share the distribution common to all the other intervals, and is not concentrated at the end of the interval, perhaps reflecting bidders' attempts to beat the deadline, and perhaps reflecting that not all bids that bidders intend to make in the last minute are successfully transmitted. In Amazon, on the other hand, averaging over all end intervals with $T > 10$ minutes in Figure 5, only about 25% of all bids are submitted in the first decile (with a mean deviation of 8%). Similarly to eBay, the last ten minutes on Amazon are less weighted towards the end than the longer intervals. This is reflected also in the first three deciles of the $T = 30$ minutes distribution (representing the last 3, 6, and 9 minutes) and by the first decile of the $T = 1$ hour distribution (representing the last 6 minutes).

Equation (1) characterizes the scale independence (i.e., independence of T) of the distributions F_T for end intervals $[0,T]$. It states that the probability that the last bid is submitted in the last at seconds given the auction lasts t more seconds only depends on the percentile (a) but not on the length of the interval (t).

(1) $$\frac{F_T(at)}{F_T(t)} =: g(a) \text{ for all } a \in [0,1] \text{ and all } 0 \le t \le T.$$

It follows for all $b \in [0,1]$ that

$$g(ab) = \frac{F_T(abt)}{F_T(t)} = \frac{F_T(abt)}{F_T(at)}\frac{F_T(at)}{F_T(t)} = g(b)g(a).$$

47 Applying χ^2-tests, pairwise comparisons of the F_T graphs reveal that these effects are statistically significant at the 5%-level for eBay and Amazon, respectively.

The solution to this functional equation (apparently first studied by Cauchy, see, e.g. Açzel, 1966, p. 41), is $g(a) = a^{\alpha}$ for some α. It follows that

$$\frac{F_T(at)}{F_T(t)} = a^{\alpha} \text{ and } F_T(at) = F_T(t)\,a^{\alpha}.$$

Since, in our application, $0 \le t \le T$, setting $t = aT$ and noting that $F_T(T) = 1$, we have that the solution to the functional equation (1) is the cumulative distribution

$$F_T(t) = \left(\frac{t}{T}\right)^{\alpha}.$$

Ordinary least-square estimates for α using as the dependent variable the logarithm of the cumulative distribution of the sample and as the explanatory variable the logarithm of t/T with T = 12 hours, yield α = 0.392 for eBay Computers (R^2= 0.99, N = 301) and 0.228 for eBay Antiques (R^2= 0.99, N = 260). The same method yields α = 0.522 for Amazon Computers (R^2= 0.97, N = 151) and 0.529 for Amazon Antiques (R^2= 0.99, N = 98). Based on these coefficient estimates, Figure 6 shows the actual and estimated cumulative distributions.[48]

[48] Note that the slope $\alpha \ge 0$ is a measure of the concentration of bids at the deadline; α = 1 implies a uniform distribution of bids over time and α = 0 implies that all last bids are submitted in the last second. Recall, however, that the timing in Amazon is computed with the help of the hypothetical deadline, which made a difference for 30 Amazon-bidders, and which makes a direct comparison harder.

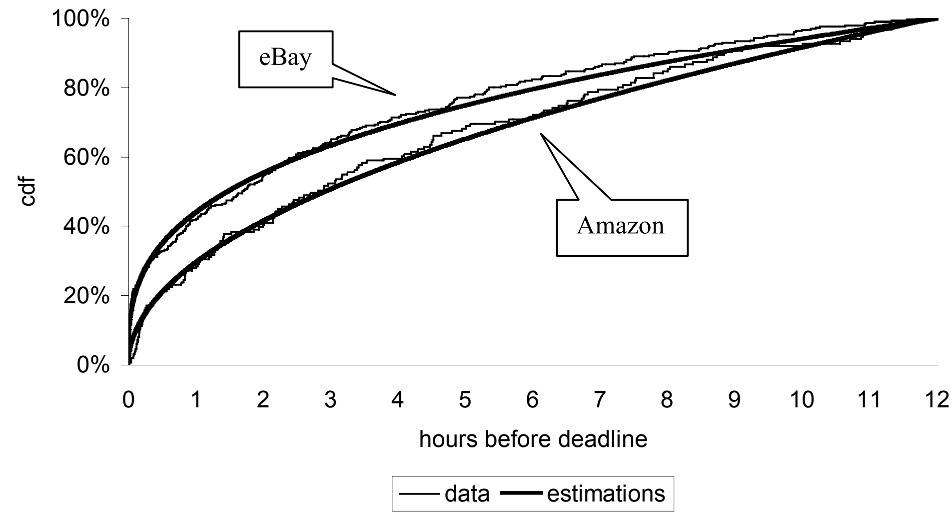

(a) Computers (eBay: α = 0.392; Amazon: α = 0.522)

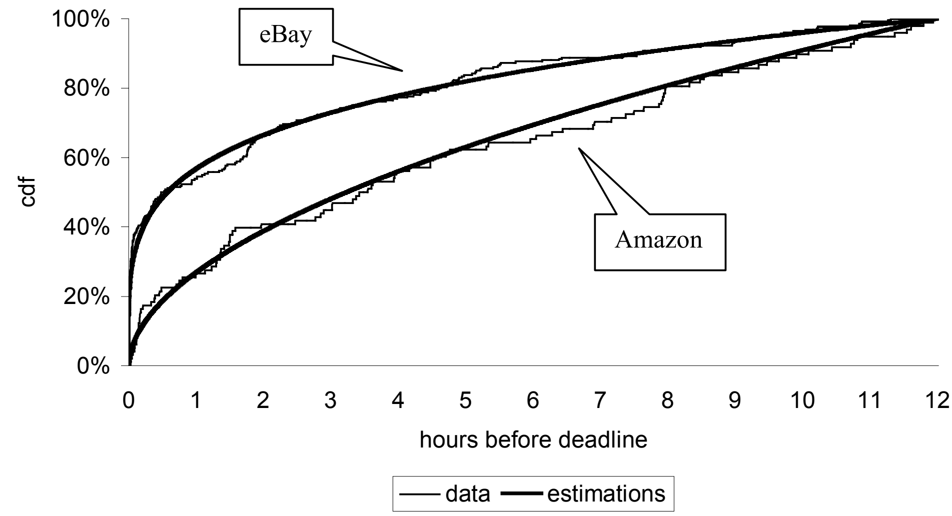

(b) Antiques (eBay: α = 0.228; Amazon: α = 0.529)

Figure 6. Estimated cumulative distribution functions $F_T(t) = (t/T)^\alpha$ and actual data for the timing of the last bids in eBay and Amazon

Scale independence of this sort has been observed with respect to many economic and social variables, yet completely satisfying explanations for this phenomenon are not always available.[49] It has also been observed in the distribution of agreements over time in bargaining experiments. As in auction bids, bargaining agreements exhibit a pronounced deadline effect; Roth, Murnighan, and Schoumaker (1988) report a striking concentration of agreements reached in the last seconds before the deadline. (Note that, as in eBay auctions, bargaining with a fixed deadline creates incentives to act very late in equilibrium *because* there is a risk of disagreement associated with waiting too long; see e.g., Fershtman and Seidmann, 1993, and Ma and Manove, 1993). Roth et al. note that all their experimental studies "show a substantial number of late agreements, even as late as the last second. Although the choice of final intervals is arbitrary, there is a clear concentration of agreements near the deadline regardless of the interval chosen to measure 'nearness' to the deadline" (p. 811). Kennan and Wilson (1993, p. 95) further analyzed the experimental data and emphasized the self-similarity that we also observe in the internet auctions:

> "(…) an increasing settlement rate is an artifact of the way time is measured. For example, Roth, Murnighan, and Schoumaker (1988, Figures 1A and 4A) reported data on 621 agreements in bargaining periods that lasted nine minutes, with 11 percent (69) of these agreements made in the first third of the bargaining period. After three minutes have passed without agreement, the remaining six minutes can be treated as a rescaling of time, with the implication that 11 percent of the remaining agreements should be made in the first third of the remaining time. In fact, the reported data show that 13 percent of the remaining 552 agreements were concluded in the first third of the last four minutes. This does not explain

---

[49] One of the first empirical regularities in economics is Pareto's law. It states that – independent of time, space, and institutions – the number of households N having income greater than x is distributed according to the Pareto distribution $N = Ax^{-\alpha}$. Most explanations for Pareto's law are based on stochastic processes involving variables such as mortality and interest rates (see, for instance, Wold and Whittle, 1957), but the income distribution may also be viewed as the result of individual decision making (see Mandelbrot, 1962; see also Persky, 1992, and the references therein for a detailed account of the history of Pareto's law). Similarly, business firm sizes, that are often found to be distributed according to Pareto distributions, can be generated by stochastic processes (Simon and Bonini, 1958). However, as Simon (1991, p. 29) puts it "Without the introduction of very particular ad hoc assumptions, unbutressed by empirical evidence, neoclassical theory provides no explanation for the repeated appearance of Pareto distributions of business firm sizes in virtually all situations where size distributions have been studied". More recently, it is found that scaling independence also occurs in the internet: the distribution of links that internet sites receive from other sites, and the distribution of visitors per site follows a power law (see Huberman et al., 1998, and Adamic and Huberman, 1999 and 2000). See also Zipf (1949) and Mandelbrot (1982) for compilations of a wide range of examples for social, economic, and other phenomena with scaling independence properties.

> why the agreement rate in the initial three minutes was 11 percent, but given any initial agreement rate, the rescaling argument explains the deadline effect."

The question why both the timing of bids in internet auctions and the timing of agreements in bargaining experiments have self-similar distributions is an open question. The empirical regularity might be deduced from stochastic models involving scaling properties of exogenous physical variables. The distribution of last bids may, for instance, reflect the distribution of bidders' availability to bid late. Furthermore, since the price increases over time, which prevents bidders from bidding late if their reservation price has been exceeded, the distribution of last bids over time might also be related to the distribution of individual valuations for the item.

VII. Discussion and Conclusions

Late bidding in internet auctions has aroused a good deal of attention. 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 source of advantage from bidding late is reduced, in comparison to the cost of delaying 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. The difference between the amount of late bidding for computers and for antiques supports the hypothesis that the bidders respond to the additional strategic incentives for late bidding in markets in which expertise plays a role in appraising values. Finally, the substantial amount of late bidding observed on Amazon, (even though substantially less than on eBay) suggests that there are also non-strategic causes of late bidding, possibly due to naivete or other non-rational cause, particularly since the evidence suggests that it is reduced with experience.

Since late bids run the risk of not being successfully transmitted, incentives to bid late may affect prices and the efficiency of allocations. But while the stakes are large,[50] little is known about how the design of the rules for ending auctions actually affect revenues and efficiency (but see the experimental study by Ariely et al., 2001). Internet auctions and other electronic markets are an especially fertile ground for future 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, 2001). Electronic markets therefore seem likely to play a major role in shaping the emerging economics of design.

References

Açzel, J. "Lectures on functional equations and their applications," New York: Academic Press, 1966.

Adamic, Lada A., and Huberman, Bernardo A. "The Nature of Markets in the World Wide Web." Working Paper, May 1999, Xerox Palo Alto Research Center.

Adamic, Lada A., and Huberman, Bernardo A. "Power-Law Distribution of the World Wide Web." *Science*, March 2000, *287*, 2115a.

Ariely, Dan, Axel Ockenfels and Alvin E. Roth. "An Experimental Analysis of Late Bidding in Internet Auctions," Working paper, Harvard University, 2001.

Bajari, Patrick and Ali Hortaçsu. "Winner's Curse, Reserve Prices and Endogenous Entry: Empirical Insights from eBay Auctions." Working paper, Stanford University, 2000.

Fershtman, Chaim and Seidmann, Daniel J. "Deadline Effects and Inefficient Delay in Bargaining with Endogenous Commitment." *Journal of Economic Theory*, August 1993, *60*(2), 306-21.

[50] For instance, eBay recently reported that the value of goods traded between the 22,500,000 registered eBay users reached $1.6 billion in the fourth quarter of 2000.

Huberman, Bernardo A., Pirolli, Peter L.T., Pitkow, James E., and Lukose, Rajan M. “Strong Regularities in World Wide Web Surfing.” *Science*, April 1998, *280*, 95-7.

Landsburg, Steven E. “My Way to the eBay: An economics professor's strategy for winning online auctions,” *Slate*, Thursday, April 8, 1999, http://slate.msn.com/Economics/99-04-08/Economics.asp.

Lucking-Reiley, David. “Vickrey Auctions in Practice: From Nineteenth-Century Philately to Twenty-First-Century E-Commerce, ” *Journal of Economic Perspectives*, 2000, 14(3), 183-192.

Ma, Ching-to Albert, and Manove, Michael. “Bargaining with Deadlines and Imperfect Player Control.” *Econometrica*, November 1993, *61*(6), 1313-39.

Malhotra, Deepak, and J. Keith Murnighan, “Milked for all Their Worth: Competitive Arousal and Escalation in the Chicago Cow Auctions,” working paper, Kellogg School of Management, Northwestern University, April, 2000.

Mandelbrot, Benoit B. “Paretian Distributions and Income Maximization,” *Quarterly Journal of Economics*, 76(1), February 1962, pp. 57-85.

Mandelbrot, Benoit B. “The fractal geometry of nature,” San Francisco: W.H. Freeman, 1982.

Milgrom, Paul. *Auction Theory for Privatization*, in preparation, 2000.

Moldovanu, Benny, and Manfred Tietzel. “Goethe’s Second Price Auction.” *Journal of Political Economy*, 1998, 106(4), 854-59.

Ockenfels, Axel and Roth, Alvin E. “The Timing of Bids in Internet Auctions: Market Design, Bidder Behavior, and Artificial Agents.” *Artificial Intelligence Magazine*, forthcoming, 2001.

Persky, Joseph. “Retrospectives: Pareto’s Law.” *Journal of Economic Perspectives*, Spring 1992, *6*(2), 181-92.

Resnick, Paul and Zeckhauser, Richard. “Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay’s Reputation System.” Draft prepared for NBER workshop, 2001.

Roth, Alvin E. “The Economist as Engineer: Game Theory, Experimental Economics and Computation as Tools of Design Economics,” Fischer Schultz lecture, *Econometrica*, forthcoming, 2001.

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*, 2001, forthcoming. (http://www.economics.harvard.edu/~aroth/papers/eBay.veryshortaer.pdf)

Simon, Herbert A. "Organizations and Markets," *The Journal of Economic Perspectives*, Spring 1991, 5(2), 25-44.

Simon, Herbert A., and Bonini, Charles P. "The Size Distribution of Business Firms." *American Economic Review*, September 1958, *48*(4), 607-17.

Vickrey, William. "Counterspeculation, Auctions, and Competitive Sealed Tenders." *Journal of Finance*, 1961, 16, 8-37.

Wilcox, Ronald T. "Experts and Amateurs: The Role of Experience in Internet Auctions," *Marketing Letters*, 2000, forthcoming.

Wold, H.O.A., and Whittle, P. "A Model Explaining the Pareto Distribution of Wealth." *Econometrica*, October 1957, *25*(4), 591-5.

Zipf, George K. "Human Behavior and the Principle of Least Effort." Cambridge: Addison-Wesley Press, 1949.