CIRJE-F-127

Strategic Implications of Uncertainty Over One's Own Private Value in Auctions

Eric B. Rasmusen
Indiana University / CIRJE

August 2001

Strategic Implications of Uncertainty Over One's Own Private Value in Auctions

August 15, 2001

Eric Rasmusen

Abstract

Bidders have to decide whether and when to incur the cost of estimating their own values in auctions. This can explain sniping– flurries of bids late in auctions with deadlines– as the result of bidders trying to avoid stimulating other bidders into examining their bid ceiling more carefully.

Visiting Professor, CIRJE, Dept. of Economics, University of Tokyo (2001), Professor of Business Economics and Public Policy and Sanjay Subhedar Faculty Fellow, Indiana University, Kelley School of Business, BU 456, 1309 E. 10th Street, Bloomington, Indiana, 47405-1701. Office: (812) 855-9219. Fax: 812-855-3354. Erasmuse@indiana.edu. Php.indiana.edu/∼erasmuse. Copies of this paper can be found at Php.indiana.edu/∼erasmuse/papers/auction.pdf. This draft has lots of algebraic steps included for checking; the published version will be more succinct.

I thank Michael Fishman, Daniel Levin, Michael Rothkopf, Jeffrey Stake, and participants at a seminar at Tokyo University for their comments. I thank Harvard Law School's Olin Center and the University of Tokyo's Center for International Research on the Japanese Economy for their hospitality.

Jeff happily awaited the end of the E-Bay auction. He'd submitted a bid ceiling of $2,100 for a custom-made analog stereo amplifier, and the highest anybody else had submitted was $1,400, so he was sure to win. Since he'd followed the advice of E-Bay and academic auction theory, submitting his true maximum price, he looked forward to a cool $700 in consumer surplus. It was five minutes before the auction deadline. And then disaster struck. The winning bid rose to $1,800, and then $1,900, and $2,000. And then it rose to $2,150, and Jeff was losing! Worse yet, as he feverishly thought hard about how much the amplifier was worth to him, he realized he actually would have been willing to pay $2,500. But by then it was too late–the auction was over.[1]

1. Introduction

In a private-value auction, the value to each bidder of the object being auctioned is independent of the value to every other bidder. The bidders still ultimately care about each other's values, since that will turn out to affect how much they have to pay to win, but other bidder's values do not convey any new information about one's own valuation. In the standard ascending auction, in fact, although a bidder would prefer that the other bidders all have low private values so he could win at a lower price, he would have no use to make of information about their values in deciding his own bidding strategy. This is in contrast to common value auctions, in which a bidder does learn something about his own value when he learns how much someone else would pay.

Auction theory begins with private-value auctions because a player's strategy is simpler; he does not have to worry about estimating the maximum amount he would be willing to pay, and the complexity is just in figuring out how to minimize the amount he pays while still winning the item. In practice, however, private-value auctions can be even more difficult for the bidder than common-value auctions. This is for an "engineering reason" absent from our

[1]From a story related to me by Jeffrey Stake. I have taken artistic license with details.

models: bidders do not costlessly know their private values. Whether we think of the problem as learning one's preferences or learning about the item being auctioned, the bidder can only discover his private value at some cost. In a business auction such as a corporate takeover, the bidder may well spend millions of dollars to learn the synergies between his own company and the target company. In a consumption auction such an estate sale, the bidder must scratch his head and agonize over whether the handsome old table will really go with his other furnishings at home. Finding out what the other bidders think would be little or no help, unlike if there is an element of common value, in which case he could hope to learn from the knowledge of some better-informed bidder as conveyed by the opening bids of the auction.

Fortunately, the bidder in an ascending private-value auction usually doesn't need to figure out the exact value of the item to himself. It is sufficient to figure out that his value is probably higher than that of the second-highest bidder in the crowd. This is perhaps why we do not think of private-value auctions as being so difficult computationally as common-value auctions. In tight competition, however, even figuring out whether one's value is the highest in the crowd is hard. And "value discovery", as I shall call it, can explain a number of odd features we observe in real-world auctions, including the story with which I started this paper, a story adapted from a real incident in the life of one of my friends. Value discovery can explain:

1. Why bidders in a private-value second-price auction would like to know how much other bidders are going to bid.
2. Why bidders update their bid ceilings in E-Bay internet auctions.
3. Why bidders use "pre-emptive bids" in ascending auctions– jumps of more than the minimum over the preceding bid.
4. Why bidders use "sniping"— the practice of submitting bids at the last minute (in a sense, the opposite of pre-emptive bids).
5. Why auction deadlines hurt the seller.
6. Why sellers prefer E-Bay auctions to ascending auctions with deadlines.

Value discovery is not the only way to explain these things. A variety of papers explain sniping, for example, based on common values (Bajari & Hortacsu [2000]), uncertainty over whether late bids will be registered by the auctioneer (Roth and Ockenfels [2001]), and irrationality (Malhotra & Murnighan [2000]). For explaining pre-emptive bids, value discovery will merely repeat in a new context the well- known explanation that entry costs- including the cost of valuing the object– can make pre-emptive bids valuable (Fishman [1988], Hirshleifer & Png [1989]). And the cost of returning to bid in an auction that takes place over several days has been shown by Carare & Rothkopf to make open-cry Dutch private auctions not equivalent to sealed-bid auctions. Value discovery is a simple idea, however, with wide application which fits certain situations particularly well and which can explain a wide range of phenomena.

The closest model to the present one is perhaps that of Dan Levin and James Smith (1994), in which bidders simultaneously decide whether to pay a certain amount to learn their private values and enter the auction. Levin and Smith calculate how many bidders will enter and compare it with what is best for efficiency and the seller. In the present paper, there will be no entry cost, but at any time during the auction a bidder has the opportunity to pay to discover his private value. I restrict the model to two bidders, only one of whom is uninformed about his value, and focus on their strategic interaction. Roland Guzman and Charles Kolstad (1999) also construct a model of a private-value auction with possible information acquisition, but since they look at a sealed-bid auction, timing is unimportant.

In another stream of research but closely related is the model of Kent Daniel and David Hirshleifer (1998). Extending the entry-cost idea of Fishman and Hirshleifer-Png, they tell a story in which each bid is costly in a private-value auction and this leads to a series of jump bids, rather than just a pre-emptive bid at the start. This is because as the auction proceeds, information is revealed, and one bidder's use of a jump bid to signal his high valuation may lead to another bidder doing the same. Their paper does have bids stimulating discontinuous behavior, as will the model here, but each player knows his own valuation, so the focus is on trying to show the other player that one's own valuation is high so as to make him give up on winning the auction.

I will proceed by laying out a model of value discovery and applying it to three sets of auction rules: second-price, ascending with deadline, and E-Bay auctions. I will begin with players who optimize nonstrategically and then go on to strategic behavior. After laying out the equilibria and showing how they illustrate the odd features mentioned above, I will discuss more implications of the model.

2. The Model

The two possible bidders in an auction, both risk-neutral, have private values v_1 and v_2 which are statistically independent.

Bidder 1 does not know his value, v_1, except that it is uniformly distributed on [0, 1]. If he wishes, at any time he can pay c and discover his value after additional time δ has passed. Bidder 2 does not observe whether Bidder 1 has paid c.

Bidder 2's value, v_2, is w with probability $1-\theta$ and $\overline{v}$ with probability θ, with $\theta \in [0,1)$, (so θ is strictly less than one), $w \in (0, .5)$, and $\overline{v} \in (.5, 1)$. He does not know v_1.

We will look at three sets of auction rules: a second-price auction, an ascending auction with a deadline, and an E-Bay auction.[2] The minimum bid increment is an arbitrarily small amount ϵ.

- In the second-price auction, each bidder submits one bid, without knowing what the other has done. Whoever submits the highest bid wins the auction, but pays the bid submitted by the other bidder, or zero if the other bidder chose not to participate.
- In the ascending auction with a deadline (hereafter "the ascending

[2] Yes, I know the company styles itself "eBay," but the rest of us need not accept debased typography.

auction"), players submit bids in continuous time until the deadline at time T. Whoever has the highest bid at time T wins the auction and pays his bid.

- In the E-Bay auction, a player submits a "bid ceiling" and the auctioneer increases the player's bid as necessary up to his ceiling until the deadline at time T. Anyone can observe the current bid posted by the auctioneer. Increasing one's ceiling costs R and may be done at any time, but reductions are not allowed. Whoever has the highest bid at time T wins and pays his bid.

Discussion of the Assumptions

Our purpose is to model a situation in which one bidder is uncertain about his value and about whether there exists another bidder who has a higher value. This bidder's decision whether to discover his value is the driving force of the model. I have here consolidated explanations of various features of the model just described.

Bidder 1 cannot discover his value instantaneously. discovery takes time as well as money. This introduces a tradeoff between discovering early, which is costly and perhaps will turn out to be wasted effort, and not discovering– since there is not time to discover late in the auction. The same tradeoff would be present if instantaneous discovery was possible but the cost of discovery rose with its speed.

It is especially important in this incomplete-information game to think of the probabilities of each type of the uninformed player as being the subjective probabilities of the uninformed player. What is most interesting is when Bidder 2 has the high value, $v_2 = \overline{v}$. The variable θ represents the strength of Bidder 1's belief that this is the actual situation. If $\theta > 0$, Bidder 1 is not completely surprised if Bidder 2 actually has a high value. If $\theta = 0$, an interesting special case, Bidder 1 puts negligible probability on $v_2 = \overline{v}$ and is surprised to learn it. The real-world application is to an auction in which Bidder 1 is a naive (but not irrational) new player and Bidder 2 is a professional on hand to take advantage of his naivete. Internet auctions are often like this.

The possible high value for Bidder 2, $v_2 = \overline{v}$, is chosen to exceed .5 so that without having discovered his own value Bidder 1 expects Bidder 2 to win the auction if v_2 takes its high value. The possible low value for Bidder 2, $v_2 = w$, is chosen to exceed 0 so that for him to signal that v_2 is high requires a jump bid over w.

The assumption that Bidder 2 does not observe Bidder 1's payment of c is made for concreteness, not to drive results. I do not know of anywhere that it would make a difference to the final outcome of the auction.

The purpose of assuming a minimum bid increment is to make determinate the strategy of a player who wants to bid the smallest amount possible that is greater than w. Otherwise, there is an open-set problem of nonexistence.

The model does not allow the seller to post a reserve price, even though he could profit by using one. If θ were large, so Bidder 2 almost certainly has $v_2 = \overline{v}$, for instance, then a reserve price of $\overline{v}$ would maximize seller profit. The usefulness of that reserve price is an artifact of the assumption that v_2 only takes two values, however, so this model is not well suited to analyzing optimal reserve prices. Its simplifications are instead designed to illuminate bidder behavior.

In many models, the E-Bay auction and the ascending auction with a deadline would be equivalent. Here, the key differences are that (a) in the E-Bay auction revising one's bid ceiling is costly and (b) in the E- Bay auction, it is not possible to publicly display "jump bids" that increase the current winning bid by more than the minimum allowed increment.

The assumption that revising one's bid in an E-Bay auction costs R, whereas the initial bid is costless and bidding in an ascending auction is costless, represents the idea that the E-Bay auction lasts several days and it is more costly for the bidder to return to his computer and update his bid than to submit an initial bid after visiting the website to see what is being auctioned. At a technical level, the impact of the assumption is that a bidder would prefer to submit one bid ceiling and never revise it, instead of being indifferent about frequent revisions.

3. Equilibrium when Bidder 2 Does Not Know that Bidder 1 Can Discover v_1

This section will serve as a benchmark and introduce the nonstrategic effects of value discovery before we add the complications of strategic reactions. It also is a realistic model in itself of the situation where bidders put trivial probability on the possibility that their bids will induce another bidder to revise his value upwards, but are wrong in the particular case at hand.[3]

3a. The Second-Price Auction

Equilibrium Behavior. Bidder 2 bids his value, either w or $\overline{v}$. Bidder 1 pays to discover his value and bid v_1 if $.5(1-\theta)w^2+.5\theta\overline{v}^2-\theta\overline{v}+.5\theta > c$, as derived below using expression (3). If he decides not to pay to discover v_1, he bids .5.

Explanation. If Bidder 1 bids .5, his expected payoff is

$$(1-\theta)(Ev_1 - w) + \theta(0) = (1-\theta)\,(.5-w) \tag{1}$$

This cannot be increased by bidding any higher or lower amount, since the expected value of the item is .5 and so Bidder 1 wishes to win only if the price he must pay, Bidder 2's bid, is .5 or less.

If Bidder 1 pays c to discover v_1 and then bids if $v_1 > w$, his expected payoff is made up of four components:

[3] It would be straightforward to construct a model in which Bidder 2 has a subjective probability γ that Bidder 1 has the possibility of paying c to discover v_1. Such a model would be even more intricate than the present one, however, and most likely the conclusion would be that if γ was low the equilibria would be similar to those here in Section 3 with $\gamma = 0$, and if γ was high the equilibria would be similar to those in Section 4 with $\gamma = 1$. This is because Bidder 2's discovery that Bidder 1 actually does have the ability to discover v_1 would come too late for Bidder 2 to act strategically in response.

1. The discovery cost, c.

2. With probability w, the auction benefit of 0, if it turns out that $v_1 < w$ and he is sure to lose the auction.

3. With probability $(\overline{v} - w)$, the benefit from bidding $v_1 < \overline{v}$, which is the benefit from winning at price w if $v_1 = w$ plus the payoff of 0 if Bidder 2 has value $\overline{v}$ and wins the auction.

4. With probability $(1 - \overline{v})$, the benefit from bidding $v_1 > \overline{v}$: the benefit of $(v_1 - w)$ if Bidder 2 bids w plus the benefit of $(v_1 - \overline{v})$ if Bidder 2 bids $\overline{v}$.

Equation (2) captures those four components.

$$
\begin{aligned}
& -c + w(0) + [\overline{v} - w]\left[(1-\theta)(E(v_1|w < v_1 < \overline{v}) - w) + \theta(0)\right] \\
& +(1-\overline{v})[(1-\theta)[E(v_1|v_1 > \overline{v}) - w] + \theta(E(v_1|v_1 > \overline{v}) - \overline{v})] \\
= \;& -c + (\overline{v} - w)(1-\theta)(.5(\overline{v} + w) - w) + [1 - \overline{v}]\left[(1-\theta)\,(.5(1+\overline{v}) - w + \theta(.5(1+\overline{v}) - \overline{v})\right] \\
= \;& -c + .5(1-\theta)\,(\overline{v}^2 - w^2) - (1-\theta)(\overline{v} - w)w - (1-\theta)(1-\overline{v})w + .5(1-\theta)(1-\overline{v})(1+\overline{v}) \\
& +.5\theta(1-\overline{v})^2 \\
= \;& -c + .5(1-\theta)\overline{v}^2 - .5(1-\theta)w^2 - (1-\theta)(1-w)w + .5(1-\theta) - .5(1-\theta)\overline{v}^2 + .5\theta(1-\overline{v})^2 \\
= \;& -c - .5(1-\theta)w^2 - (1-\theta)\,w + (1-\theta)w^2 + .5(1-\theta) + .5\theta(1-\overline{v})^2 \\
= \;& -c + .5(1-\theta)(1-w)^2 + .5\theta(1-\overline{v})^2.
\end{aligned}
\tag{2}
$$

Bidder 1 will discover information if that makes his payoff bigger than the payoff if he did not discover (expression (1), which is the case if expression

(3) is positive:

$$
\begin{aligned}
& [-c + .5(1-\theta)(1-w)^2 + .5\theta(1-\overline{v})^2] - [(1-\theta)\,(.5-w)] \\
= \;& -c + .5(1-\theta) + .5(1-\theta)w^2 - (1-\theta)w + .5\theta(1-\overline{v})^2 - .5(1-\theta) + (1-\theta)w \\
= \;& -c + .5(1-\theta)w^2 + .5\theta(1-\overline{v})^2
\end{aligned} \tag{3}
$$

Bidder 1 is more likely to decide to discover the value if c is smaller, since that makes expression (3) larger. Since its derivative with respect to $\overline{v}$ is $-\theta\overline{v}$, Bidder 1 is less likely to decide to discover the value if $\overline{v}$ is larger. The derivative with respect to θ has variable sign, depending on the relative sizes of $\overline{v}$ and w.

Both decisions–discovering and not discovering– are possible. Bidder 1 would clearly not want to pay to discover his value if c were too high. He would be willing to discover his value if c were low enough, because the remaining terms of expression (3) have a positive sum. Intuitively, he wants to pay to discover his value if c is low enough because he might find that $v_1 > \overline{v}$, in which case he would bid higher than .5 and win the auction on some occasions when he would otherwise lose it.

If a bidder knows other bidders' values he can better make the decision of whether to pay c to discover his own value. Thus, this auction demonstrates:

RESULT 1: Bidders in a private-value second-price auction may find it useful to know the values of the other bidders.

3b. The Ascending Auction

Equilibrium Behavior. Bidder 2 bids up to v_2 as necessary.

Suppose $c < .5(1-\overline{v})^2$. Bidder 1 will pay to discover his value at the beginning of the auction if $c < .5w^2$, and then bid up to v_1 as necessary. Otherwise, he will bid up to w and then pay to discover his value if Bidder 2 bids higher, and then use bid up to v_1 as necessary.

Suppose $c > .5(1-\overline{v})^2$. Bidder 1 will pay to discover his value at the beginning of the auction if $c < .5((1-\theta)w^2 + \theta(1-\overline{v})^2)$, and then bid up to v_1 as necessary. Otherwise, he will bid up to .5 and then stop.

Explanation. Bidder 2's strategy is straightforward. Given that he does not know that Bidder 1 can acquire information, Bidder 1 thinks this is a standard ascending auction, so he simply bids up to his reservation value.

Bidder 1 has to make the choice of if and when to acquire information about v_1. If he does not acquire information, he should simply bid up to .5, his best estimate of v_1.

Let us work back from the end. Suppose Bidder 1 observes Bidder 2 bid $w + \epsilon$. Bidder 1 deduces that $v_2 = \overline{v}$ and he will lose the auction unless he acquires more information, for a payoff of zero. If he does pay for value discovery his expected payoff viewed at that moment is

$$\begin{aligned} & -c + Prob(v_1 > \overline{v})[E(v_1|v_1 > \overline{v}) - \overline{v}] \\ = \; & -c + (1-\overline{v})\left(.5(1+\overline{v}) - \overline{v}\right) \qquad (4) \\ = \; & -c + .5(1-\overline{v})^2 \end{aligned}$$

If $v_2 = \overline{v}$, Bidder 1 thus pays to discover v_1 if $c < .5(1-\overline{v})^2$. His payoff viewed from the start of the game if $c > .5(1-\overline{v})^2$, so he does not pay to discover v_1 after finding out that $v_2 = \overline{v}$ is

$$(1-\theta)(.5-w) \qquad (5)$$

If, on the other hand, $c < .5(1-\overline{v})^2$, so he does pay to discover v_1 after finding out that $v_2 = \overline{v}$, his expected payoff viewed from the start of the game is

$$(1-\theta)(.5-w) + \theta(-c + .5(1-\overline{v})^2) \qquad (6)$$

Another alternative, however, is for Bidder 1 to pay to discover v_1 at the very beginning of the auction. The advantage of that is that if $v_1 < w$ he can avoid winning with a bid of w.

Bidder 1's expected payoff from acquiring information at the start of the auction is

$$
\begin{aligned}
& -c + [1-\theta][Prob(v_1 < w)[0] + Prob(v_1 > w)[E(v_1|v_1 > w) - w]] + \\
& \theta(Prob(v_1 < \overline{v})[0] + Prob(v_1 > \overline{v})[E(v_1|v_1 > \overline{v}) - \overline{v}]) \\
= \;& -c + [1-\theta][(1-w)(.5(1+w) - w)] + \theta((1-\overline{v}).5(1+\overline{v}) - \overline{v}]) \\
= \;& -c + [1-\theta](.5(1-w)^2)) + \theta((.5(1-\overline{v})^2))
\end{aligned}
\tag{7}
$$

Suppose first that $c < .5(1-\overline{v})^2$, so the best alternative to paying c at the start of the auction is to pay it upon discovering that $v_2 = \overline{v}$, which has overall payoff (6). Payoff (7's advantage over payoff (6) is

$$
\begin{aligned}
& -c + [1-\theta](.5(1-w)^2)) + \theta((.5(1-\overline{v})^2)) - [(1-\theta)(.5-w) + \theta(-c + .5(1-\overline{v})^2)] \\
= \;& -c + [1-\theta](.5(1-2w+w^2) - (.5-w)) - \theta c \\
= \;& (1-\theta)(.5w^2 - c)
\end{aligned}
\tag{8}
$$

Acquiring information at the start of the auction is therefore best if $c < .5(1-\overline{v})^2$ and $c < .5w^2$.

Suppose next that $c > .5(1-\overline{v})^2$, so the best alternative to paying c at the start of the auction is never to pay it, which has overall payoff (5). Payoff (7's advantage over payoff (5) is

$$
\begin{aligned}
& -c + [1-\theta](.5(1-w)^2)) + \theta((.5(1-\overline{v})^2)) - (1-\theta)(.5-w) \\
= \;& -c + [1-\theta](.5(1-2w+w^2) - (.5-w)) + \theta((.5(1-\overline{v})^2)) \qquad (9) \\
= \;& -c + .5(1-\theta)w^2 + .5\theta(1-\overline{v})^2
\end{aligned}
$$

Acquiring information at the start of the auction is therefore best if $c > .5(1-\overline{v})^2$ but $c < .5((1-\theta)w^2 + \theta(1-\overline{v})^2)$.

Thus, if w is large and c is small, Bidder 1 is more likely to wish to pay to discover his value. As one might expect, if c were 0, Bidder 1 would pay to discover his value regardless of the value of w.

Here in Section 3, Bidder 2 does not know that Bidder 1 can discover v_1. Therefore, Bidder 2 will be shocked if he bids .5 and Bidder 1 continues to bid higher; Bidder 2 would have expected that Bidder 1 would deduce that $v_1 = \overline{v}$, greater than Bidder 1's expected value of $v_1 = .5$. Nonetheless, Bidder 2 has a dominant strategy— keep matching the bid until the winning bid exceeds $v_2 = \overline{v}$— and a dominant strategy is a best reply even to apparently irrational behavior by another player.

3c. The E-Bay Auction

Equilibrium Behavior. Bidder 2 submits a maximum bid of v_2: either w or $\overline{v}$.

Bidder 1 must decide whether to discover his value before learning anything about v_2, after learning that $v_2 = \overline{v}$, or never. If

$$-(1-\theta)(.5-w) + .5\theta(1-\overline{v})^2 + Max([-c + .5(1-\theta)w], [-\theta c]) < 0. \quad (10)$$

then he submits a bid ceiling of .5 and never pays to discover his value (see equation (16) below).

If inequality (10) is false, Bidder 1 discovers his value even before learning anything about v_2 and submits a bid ceiling of v_1 if (see equation (14) below)

$$(1-\theta)c < .5(1-\theta)w \quad (11)$$

If both (10) and (11) are false, Bidder 1 first submits a bid ceiling of .5. If the posted maximum bid rises to .5, he pays c to discover his value, and pays R to submit a new bid ceiling of v_1 if he discovers that $v_1 > \overline{v} + R$.

Explanation. Now, if Bidder 1 decides to revise his bid after seeing what Bidder 2 does, he must pay an additional R. This will tend to discourage his acquiring information after the start of the auction and encourage acquiring it before the auction starts.

Bidder 1's payoff from the strategy of discovering v_1 without learning anything about v_2 is

$$
\begin{aligned}
& -c + Prob(v_1 < w)(0) + Prob(v_1 > w)(1-\theta)(E(v_1|v_1 > w) - w) \\
& +\theta[Prob(w < v_1 < \overline{v})(0) + Prob(v_1 > \overline{v})(E(v_1|v_1 > \overline{v}) - \overline{v})] \\
= \;& -c + 0 + (1-\theta)[.5(1+w) - w] + \theta[0 + (1-\overline{v})(.5(1+\overline{v}) - \overline{v})] \\
= \;& -c + .5(1-\theta) + .5(1-\theta)w - (1-\theta)w + .5\theta(1-\overline{v})(1-\overline{v}) \\
= \;& -c + .5(1-\theta) + .5(1-\theta)w - (1-\theta)w + .5\theta(1-\overline{v})^2 \\
= \;& -c + .5(1-\theta) + .5(1-\theta)w - (1-\theta)w + .5\theta(1-\overline{v})^2 \\
= \;& (1-\theta)(.5-w) - c + .5(1-\theta)w + .5\theta(1-\overline{v})^2.
\end{aligned}
\tag{12}
$$

Suppose Bidder 1 did not discover his value. The price is bid up to .5 if $v_2 = \overline{v}$. At that point, Bidder 1 will decide either to drop out or to pay c and discover his value, after which, if $v_1 > \overline{v} + R$ he would pay R to increase his bid ceiling. Bidder 1's payoff for the entire game from the strategy of paying c to discover his value only if he deduces that $v_2 = \overline{v}$ is

$$
\begin{aligned}
& (1-\theta)[Ev_1 - w] + \theta[-c + Prob(v_1 < \overline{v} + R)[0] \\
& +Prob(v_1 > \overline{v} + R)[-R + E(v_1|v_1 > \overline{v} + R) - \overline{v}] \\
= \;& [(1-\theta)(.5-w)] - \theta c + \theta[0 + (1-\overline{v}-R)(-R + .5(1+\overline{v}+R) - \overline{v})] \\
= \;& [(1-\theta)(.5-w)] - \theta c + \theta(1-\overline{v}-R)(-.5R + .5(1-\overline{v})) \\
= \;& [(1-\theta)(.5-w)] - \theta c - .5\theta(1-\overline{v})R - .5\theta R^2 + .5\theta(1-\overline{v})^2.
\end{aligned}
\tag{13}
$$

The advantage of early discovery over late is therefore

$$
[-c+.5(1-\theta)w]-[-\theta c-.5\theta(1-\overline{v})R-.5\theta R^2] = -(1-\theta)c+.5(1-\theta)w+.5\theta(1-\overline{v})R+.5\theta R^2 \tag{14}
$$

as is evident from comparing the last lines of equations (12) and (13). Late discovery has the advantage that delay may reveal that $v_2 = \overline{v}$, in which case paying c to discover v_1 is unnecessary. Early discovery is better if c and $\overline{v}$ are smaller and and w and R are larger.

Bidder 1's payoff from the strategy of never discovering his value is

$$(1-\theta)(Ev_1 - w) + \theta(0) = (1-\theta)(.5 - w) \tag{15}$$

Never discovering is superior to either early or late discovery if

$$-.5\theta(1-\overline{v})^2 - Max([-c + .5(1-\theta)w], [-\theta c - .5\theta(1-\overline{v})R - .5\theta R^2]) > 0 \tag{16}$$

Each of the three strategies can be an equilibrium strategy, depending on the parameters. Never discovering is best if c is large enough, which makes expression (16) negative. Early discovery is best if c is small enough and R is large enough, in which case inequality (16) is false and expression (14) is positive. Late discovery is best if R and w are small enough, which makes inequality (16) false and expression (14) negative.

This auction demonstrates:

RESULT 2: In an E-Bay auction we may see a bidder increasing his bid ceiling.

Without value discovery, a bidder's dominant strategy is to submit his true value as his bid ceiling, rather than wait and incur extra cost R to revise his ceiling. In the present model, that is what would happen if either Bidder 1 either knew v_1 at the start or if he could not spend c to discover v_1. Here, however, a bidder can find it optimal to delay in the hope of not having to pay the cost of value discovery, but then to pay it and revise his bid ceiling if he finds there is close bidding competition.

Increases in a bidder's ceiling will not occur in Section 4, where Bidder 2 will know that Bidder 1 can discover v_1, because Bidder 2 will have incentive to delay his bid so as not to give Bidder 1 time to discover v_1.

Thinking of a bidder's estimate of his private value as changing during the course of an auction can also lead to Result 2 in another way, not present in this model, via exogenous changes in a bidder's value, discovered by him immediately at zero cost. A rational bidder will take into account the possibility of random shocks in his value when he first submits his bid ceiling, and it provides an incentive to submit the ceiling as late as possible (thus providing an explanation, though an unrealistic one, for "sniping"). If, however, there is an extra cost to submitting one's first bid ceiling late instead of early (because the bidder is already at the website early in the game but must make a special trip to return later, for example), a bidder might take the risk of submitting early. This could explain not only upward revisions in bid ceilings, but downward ones, if they were permitted. As the *Wall Street Journal* explains:

> "... on all the major sites you must submit your maximum bid up front, and you're obliged to pay it if bidding gets that high and doesn't go higher, even if you've lost interest in an item during the course of the auction."[4]

4. Equilibrium when Bidder 2 is Aware that Bidder 1 Can Pay c to Discover v_1

Let us now examine the equilibrium when Bidder 2 is aware that Bidder 1 can pay c to discover v_1.

4a. The Second-Price Auction

Equilibrium Behavior. Bidder 2 bids his value, either w or $\overline{v}$. Bidder 1 pays to discover his value and bids v_1 if $.5(1-\theta)w^2 + .5\theta\overline{v}^2 - \theta\overline{v} + .5\theta > c$, as derived earlier in inequality (3). If he decides not to pay to discover v_1 he bids .5.

Explanation. This auction has the same equilibrium as in Section 3a. Bidder 2 now knows Bidder 1 has discovered his value before the auction, for some parameters. That makes no difference, however, and the equilibrium is the

[4] "As eBay Rivals Emerge, Some Tips on Bidding And Selling on the Web," *Wall Street Journal*, Bart Ziegler, September 23, 1999, p. B1.

same as in Section 3a since submitting his true value is a dominant strategy for a bidder in this auction, in which no information comes out in the course of the bidding.

4b. The Ascending Auction

Equilibrium Behavior. Let us distinguish two cases.

Case 1 (similar to Section 3b). $c > .5(1 - \overline{v})^2$.

If $v_2 = w$, Bidder 2 stands willing to bid up to w as necessary. It does not matter when he bids.

If $v_2 = \overline{v}$, then Bidder 2 makes a pre-emptive bid of $w + \epsilon$ at the start of the auction and bids up to $\overline{v}$ as necessary.

Bidder 1 waits to see whether Bidder 2 bids $w + \epsilon$ at the start of the auction. If Bidder 2 bids $w + \epsilon$, Bidder 1 submits a bid ceiling of .5, but expects to lose. If Bidder 2 does not bid $w + \epsilon$, Bidder 1 deduces that $v_2 = w$. He then pays to discover v_1 and bids up to it as necessary if $c < .5(1 - \theta)w^2 + .5\theta(1 - \overline{v})^2$. If he does not pay to discover v_1, he bids up to .5 as necessary.

Case 2. $c < .5(1 - \overline{v})^2$.

Bidder 2 bids no more than w until time $T - \delta$, and then bids up to v_2 as necessary.

Bidder 1 pays to discover v_1 and bids up to it as necessary if $c < .5(1 - \theta)w^2 + .5\theta(1 - \overline{v})^2$. Otherwise, he bids up to .5 as necessary.

If, out of equilibrium, Bidder 2 bids more than .5 before $T - \delta$, then Bidder 1 pays to discover v_1 if he has not discovered it already. After discovering v_1, he bids up to it as necessary. Any out-of-equilibrium belief by Bidder 1 about the probability that $v_2 = \overline{v}$ supports this equilibrium.

Explanation. What has changed from Section 3b is that now Bidder 2 is aware that his bid may stimulate Bidder 1 to pay to discover v_1. Bidder 2 will use a pre-emptive bid if it will deter Bidder 1, which is what happens under the parameters of Case 1. In Case 2, however, a pre-emptive bid would fail to deter Bidder 1, and might even incite him to discover his value when he would not do so otherwise.

The parameter values are taken from the analysis of Sections 3a and 3b, where expression (4) showed that if $c < .5(1-\overline{v})^2$ Bidder 1 is willing to pay to discover his value after learning that $v_2 = \overline{v}$, and expression (9) showed that if $c < .5(1-\theta)w^2 + .5\theta(1-\overline{v})^2$, Bidder 1 is willing to pay to discover his value without learning anything about v_2.

Since a bid of more than w before $T-\delta$ is out-of-equilibrium behavior, Bidder 1 cannot use Bayes's Rule to update his belief about the probability that $v_2 = \overline{v}$. Bidding more than w when $v_2 = w$ is irrational behavior, to be sure, but so is bidding more than w when $v_2 = \overline{v}$. It turns out, however, that Bidder 1's out-of-equilibrium belief about v_2 does not matter to his optimal strategy of paying to discover v_1. If his belief is that a deviation by Bidder 2 shows that $v_2 = \overline{v}$ with certainty, then in Case 2 Bidder 1 is willing to pay to discover v_1, as we know having used expression (4) to define Case 2 on this basis. If his belief is that a deviation by Bidder 2 shows any lower probability of $v_2 = \overline{v}$, then Bidder 1 is *a fortiori* willing to pay to discover v_1, since if it turns out that $v_1 > .5 + R$ he will be able to win at the lower price of .5 instead of $\overline{v}$.

This auction demonstrates

RESULT 3: In an ascending auction we may see a bidder using a pre-emptive bid— a jump of more than the minimum required over the preceding bid.

Pre-emptive bids are a difference between sealed-bid second-price and ascending auctions. In a sealed-bid auction, pre-emptive bids are not possible, because only the bidder knows his bid. Curiously, the result of discovery cost is that the price paid by the winner in the ascending auction is no longer close to either the loser's bid or the loser's private value. It is greater than the loser's bid— if we can call a non-existent bid zero— and less than the loser's expected value— because if it were not, the pre-emptive bid would fail to work.

Pre-emptive bids help both bidders. Bidder 1 is helped because he could forestall a pre-emptive bid by discovering his value before the auction begins, but he prefers not to– he would rather save the discovery cost if $v_2 = \overline{v}$.

Bidder 2 is helped because the pre-emptive bid prevents Bidder 1 from discovering v_1. There would be some cases where $v_1 > \overline{v}$, yet Bidder 1 would bid .5 and lose because he has not discovered that fact.

In this model there are no entry costs, which are another way to explain pre-emptive bids. Bidder 1 enters and bids .5 even if he is deterred from discovering v_1 by a pre-emptive bid. If there were an entry cost, then Bidder 2 would have an additional reason to use a pre- emptive bid. If there were an entry cost, a pre-emptive bid would deter Bidder 1 from entering at all, and Bidder 2 could win at a price lower than .5.

The possibility of discovery, however, can mean that Bidder 2 might choose not to use a pre-emptive bid that he would otherwise use. The reason is that Bidder 1 could respond to learning that v_2 is high in two ways: by giving up and just bidding .5, or by paying to discover his value and then bidding it and winning if v_1 is high enough. If it is the second option that is most attractive, Bidder 2 will not wish to bid early. Instead he will use the sniping strategy described above, leading us to Result 4.

RESULT 4: "Sniping" can occur in equilibrium. A bidder may purposely delay overbidding the current bid until near the auction deadline.

The combined effects of sniping and pre-emptive bids gives us Result 5.

RESULT 5: Auction deadlines hurt the seller.

In both, the reason is the same: a bidder with uncertain value decides not to try to improve his information. Sniping and pre-emptive bids both help reduce the winning bid and thus hurt the seller. Sniping is possible only becauase of the deadline T. We will see Result 5 reappear in the Section 4c in the E-Bay auction, since it, too, has a deadline. Since the result has already appeared, however, it is perhaps appropriate to discuss it here. The question is why E-Bay uses a deadline. E-Bay itself is neither buyer nor seller, selling auction services rather than the objects themselves, but it needs to attract sellers.

The first reason is the practical one that without some deadline, the auction would never end. Even an ascending auction without a deadline has a deadline of sorts: the "going, going, gone," an interval like in the after which if no new bids have been made, the auction ends.

This reason is undermined, however, by the possibility of using other ways to end the auction in finite time.

Amazon.com internet auctions can last several days, like the E-Bay auctions, but they end more like ascending auctions. The deadline is not strict. Instead, the auction ends only when (a) the deadline has passed, and (b) there has been no bidding for a certain time interval. Only if it takes more than that interval for value discovery would we see last-minute flurries of bids and bidders who would like the auction to stay open but who dare not bid higher immediately. Whether bidding takes more than that number of minutes depends on the particular bidder and object.

Another possibility, not yet employed in real auctions to my knowledge, is to use surprise endings. The rule could be that the auction lasts at least T minutes, after which it ends each minute with fixed probability Z. Since Bidder 2 has some probability of winning at a price that yields him surplus even if his bid provokes Bidder 1 to discover v_1, but has no probability of winning if he does not bid before the auction ends, a sufficiently low Z would make Bidder 2 prefer to bid early while still allowing the auction to end in finite time.[5]

[5] On the other hand, use of a surprise ending like this would, I conjecture, exacerbate the possibility that bids are low for a different reason. Roth and Ockenfels (2000) note that if there is some chance an attempted bid near a deadline will not be registered, bidders will bid low early in the hopes that all the late bids will fail to register, and will refrain from overbidding early because they know it will just stimulate a bidding war. If the ending time is stochastic, late bids might not be registered, so the rules, rather than technology,

The second reason why E-Bay might wish to use a deadline is more robust. If it is important to the seller to attract bidders, he will do things that transfer surplus from himself to bidders— here, a deadline to encourage sniping. The benefit to the sellers is less clear in this case, however, since sniping hurts bidders who have the possibility of value discovery, either by depriving them of the time to do it or by requiring them to incur the cost early as a precaution. Thus, it is useful to the seller only if he is concerned more with attracting sniping bidders than driving away bidders whose private values are uncertain.

4c. The E-Bay Auction

Equilibrium Behavior. Bidder 2 waits until $T-\delta$ to submit a bid ceiling. If $v_2 = \overline{v}$, he submits a ceiling of $\overline{v}$. Otherwise he submits a ceiling of $v_2 = w$.

Bidder 1 pays to discover v_1 and submits it as his bid ceiling if $.5(1-\theta)w^2 + .5\theta\overline{v}^2 - \theta\overline{v} + .5\theta > c$. If he decides not to discover v_1 he submits a bid ceiling of .5.

If, out of equilibrium, Bidder 1 sees the current posted bid rise to more than .5 before $T-\delta$, he pays to discover v_1 if he has not discovered it already. After discovering v_1, he bids up to it as necessary. Any out-of-equilibrium belief by Bidder 1 about the probability that $v_2 = \overline{v}$ supports this equilibrium. indicating a bid by Bidder 2, he discovers his value if $c < .5(1-\overline{v})^2$.

Bidder 2 bids no more than w until time $T-\delta$, and then bids up to v_2 as necessary.

Bidder 1 pays to discover v_1 and bids up to it as necessary if $c < .5(1-\theta)w^2 + .5\theta(1-\overline{v})^2$. Otherwise, he bids up to .5 as necessary.

Explanation. In the E-Bay auction, a pre-emptive bid cannot be used. Sniping, however, is still possible. Thus, Bidder 1 uses a strategy similar to his second-price auction strategy and Bidder 2 uses a strategy similar to the sniping part of his ascending auction strategy.

---

might lead to their phenomenon.

If Bidder 1 does not discover v_1 his payoff is

$$(1-\theta)\left(Ev_1 - w\right) + \theta\left(0\right) = (1-\theta)\left(.5 - w\right). \tag{17}$$

If Bidder 1 does discover v_1 without learning anything about v_2, his payoff is

$$\begin{array}{l} -c + Prob(v_1 < w)(0) + (1-\theta)(E(v_1|v_1 > w) - w) \\ +\theta[Prob(w < v_1 < \overline{v})(0) + Prob(v_1 > \overline{v})(Ev_1|v_1 > \overline{v}) - \overline{v})] \end{array} \tag{18}$$

This is the same expression as for the second-price auction, equation (2). Hence, we get the same condition for Bidder 1 to acquire information as we did in Section 3c.

Bidder 2's strategy of delay will yield him an expected payoff of $\overline{v} - .5$, if $v_2 = \overline{v}$ and Bidder 1 has not acquired information, If, on the other hand, Bidder 2 did not delay and Bidder 1 were stimulated to discover information, Bidder 2's expected payoff if $v_2 = \overline{v}$ would be

$$\begin{array}{rl} & Prob.\left(v_1 < .5\right)\left(\overline{v} - .5\right) + Prob.\left(.5 < v_1 < \overline{v}\right)\left(\overline{v} - E\left(v_1 \middle| \left(.5 < v_1 < \overline{v}\right)\right)\right) \\ & +Prob.\left(v_1 > \overline{v}\right)\left(0\right) \\ = & .5\left(\overline{v} - .5\right) + \left(\overline{v} - .5\right)\left(\overline{v} - .5(.5 + \overline{v})\right) \end{array} \tag{19}$$

The delay payoff, $\overline{v} - .5$, is greater than the non-delay payoff, (19), because $\overline{v} > .5$

Bidder 2 gains by delay if that prevents Bidder 1 from discovering v_1. If Bidder 1 has discovered v_1 or would never try to discover it in any case, or if $v_2 = w$, Bidder 2 might as well follow the same strategy, since delay neither helps nor hurts him.

The behavior by Bidder 1 in response to out-of-equilibrium bids by Bidder 2 is the same as in Section 4b and for the same reasons.

This auction demonstrates Result 6.

RESULT 6: The seller prefers the E-Bay auction to the ascending auction with a deadline.

Pre-emptive bids benefit both bidders because they reduce transaction costs, but they hurt the seller, because they reduce the winning bid. Hence, sellers should prefer auction rules such as the E-Bay rules which do not allow jump bids. An open-cry ascending auction could, of course, also be structured that way.[6]

Comparing Explanations. Alvin Roth and Axel Ockenfels (2000) have a different explanation for sniping. Their model is driven by the possibility that players making bids in the last minute may find the computer has not been able to get their bids in time. In that case, players will submit low bids early in the auction and higher bids in the last minute. There is some chance that none of the high bids will be accepted, and so some bidder wins with his very low initial bid. If, on the other hand, someone tries bidding higher before the last minute, so his bid definitely reaches E-Bay, he finds that other bidders will outbid him and the resulting bidding war will leave even the winner with a low payoff.[7]

Or, it could simply be that the auction is not a private-value auction at all, but a common-value auction, in which case sniping can easily arise if there is time required for updating valuations. Since the value is common to all bidders, whenever a bidder raises his bid, the other bidders will revise their value estimates upwards and bid more, to his detriment. Hence, he will delay bidding until it is too late for them to revise their estimates.

Revising a common value, however, though mathematically more com-

[6] A caveat to this seller preference against pre-emptive bids is that if transaction costs are extremely high the seller may wish to allow them so as to encourage the first bidder to pay a cost of entering the auction, something which does not arise in the present model because there is no entry cost.

[7] This has some similarity to Bertrand models of price competition, where marginal-cost pricing is Nash, but weakly dominated, and so disappears when noise is added. See, for example, Maarten Janssen & Eric Rasmusen (2001), in which N identical firms each may be active or inactive, and post prices for a consumer. Each firm knows that it might be the only active firm, so it charges a price higher than marginal cost, using a mixed strategy. As in Roth & Ockenfels, a little bit of noise in a Bertrand model (in their case, the possibility last-minute bids might not get through) results in the competing price setters ending up with positive expected profits.

plicated than revising a private value, is much more straightforward precisely because it is pure mathematics. No searching of one's preferences or spreadsheet analysis of business conditions is needed. Moreover, the upward direction of the revision, at least, is clear, unlike in private value revision. That is perhaps the key difference from the value discovery in the present paper: high bids in a common- value auction are good news (and lack of them bad news), rather than being neutral about one's valuation.[8]

5. Concluding Remarks

Value discovery explains the flurry of last-minute bidding in internet auctions as being the result of bidder ignorance of their private values. Some bidders bid late so as to prevent other bidders from having time to acquire more precise information on how much they value the object being auctioned. This also explains repeat bidding: bidders refrain from incurring the cost of thinking hard about their values until they see that bidding is high enough that such thinking is necessary.

There is a curious nonmonotonicity in the willingness of the bidder to pay to discover his private value. If he believes he is almost certain to lose the auction, he will not bother to discover his value. But if he believes he is almost certain to win the auction, he also will not bother to discover his value— for in that case, all that matters is that his bid exceed that of the second-highest bidder. In between, however, where the bidder is uncertain about whether he will win or not, it becomes useful to know his value precisely.

Value discovery may well have useful application to other kinds of markets. The auction story parallels a bidder's decision when buying at a posted price. If he knows that the object's price is much more than its value to him, he will not agonize over how much it is worth to him, and similarly if the

[8]It would also be possible to have value discovery on top of Bayesian common value updating in a common- value auction. It would take the form of paying a cost to obtain a better signal of the common value. This would provide an additional reason not to bid early in a common-value auction.

price is far less than its value. Only when the price is close to the estimate value does it become worthwhile to spend time and energy improving the estimate. The only difference is that in an auction the buyer must decide whether to do his improvement in advance, for fear that the final price will be closer to his estimate than the present one.[9]

For what kind of auctions is this model reasonable? Certainly it applies to ascending auctions conducted over a long period of time (e.g., three days), like the E-Bay auctions. It also applies to sequential auctions, like the FCC Spectrum Auction, in which sealed bids are submitted, the winning bid is announced, and then other rounds are held till nobody wants to bid higher. But it even applies to classic English auctions like those at estate sales. The auction only lasts five minutes, but bidders must decide beforehand whether to learn the value, and if they see bidding is low at first, they can spend a minute learning more about their private value and it will not be too late to enter the auction.

The model has assumed that value discovery is useful only insofar as it helps the bidder know how much to bid in the auction. This is generally true for private-value auctions of consumption goods, where the bidder does not really care about the size of his consumer surplus after winning the auction, so long as it is non-negative.

Private-value auctions of production goods are often different. There, the private value is the profit the particular bidder can make from the object being auctioned, which depends on its interaction with his other assets. The process of finding the value of this interaction usually increases the value itself, by helping the bidder better understand what to do with his assets if he wins (it can even be beneficial if he loses, since he learns something about his existing assets). In the model of this paper, this might take the form of

[9]The idea that finding out one's own value for an object explains odd behavior also shows up in bargaining. In Rasmusen (2000), I model negotiation as a process in which one player makes offers whose value the other player can determine only at some cost. This usually results in a mixed-strategy equilibrium in which the offers are sometimes high and sometimes low value, and the ignorant player sometimes investigates before accepting and sometimes does not.

increasing v_1 by some amount d if Bidder 1 pays c to discover v_1.

This distinction between consumption-good auctions and production-good auctions is not hard-and-fast. Consumption goods, too, have synergy with existing consumption goods, synergy which can be increased by careful planning. Examples of production goods for which value discovery has no post- auction value are harder to imagine, but one example is when I am buying the object for resale by a second auction in a market I monopolize. I will get the same price in the second auction regardless of whether I know that price in advance, and so my value discovery in the first auction has no post-auction value.

References

Bajari, Patrick & Ali Hortacsu (2000) "Winner's Curse, Reserve Prices and Endogenous Entry: Empirical Insights from eBay Auctions," Stanford working paper, 2000.

Carare, Octavian & Michael Rothkopf (2001) "Slow Dutch Auctions," Rutcor Research Report RRR42-2001, Rutgers University, http://rutcor/rutgers.edu/∼rrr.

Daniel, Kent & David Hirshleifer (1998) "A Theory of Costly Sequential Bidding," working paper, Northwestern University, July 30, 1998, http://kent.kellogg.nwu.edu/papers/DH.pdf (July 23, 2001).

Fishman, Michael J. (1988) "A Theory of Preemptive Takeover Bidding," *Rand Journal of Economics,* 19:88-101 (Spring 1988).

Guzman, Rolando & Charles Kolstad (1999) "Auction Equilibrium with Costly Information Acquisition," University of California-Santa Barbara Dept. of Economics working paper, May 1999.

Hirshleifer, David & Ivan Png (1989) "Facilitation of Competing Bids and the Price of a Takeover Target," *Review of Financial Studies*, 2:587-606.

Janssen, Maarten & Eric Rasmusen (2000) "Bertrand Competition Under Uncertainty," forthcoming, *Journal of Industrial Economics*, http://Php.indiana.edu/∼erasmuse/papers/bertrand.pdf.

Landsburg, Steven (1999) "My Way to the eBay: An Economics Professor's Strategy for Winning Online Auctions," *Slate*, April 8, 1999, http://slate.msn.com/Economics/99-04-08/Economics.asp.

Levin, Dan & James L. Smith, (1994) "Equilibrium in Auctions with Entry," *American Economic Review,* 84: 585-99 (June 1994).

Malhotra, Deepak & J. Keith Murnighan (2000) "Milked for All Their Worth: Competitive Arousal and Escalation in the Chicago Cow Auctions," working paper, Kellogg School of Management, Northwestern University, April 2000.

Rasmusen, Eric (2001) "Explaining Incomplete Contracts as the Result of Contract- Reading Costs," http://Php.indiana.edu/∼erasmuse/papers/negot.pdf.

Rasmusen, Eric (2001) *Games and Information*, 3rd edition. Oxford: Blackwell Publishers.

Roth, Alvin E. & Axel Ockenfels (2000) "Last Minute Bidding And The Rules For Ending Second-Price Auctions: Theory And Evidence From A Natural Experiment On The Internet," NBER Working Paper 7729, June 2000, http://www.nber.org/papers/w7729.

Wilcox, Ronald (2000) "Experts and Amateurs: The Role of Experience in Internet Auctions," *Marketing Letters*, 11: 363-374 (November 2000).