

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

First draft: March 2001

This draft: January 26, 2003

Eric Rasmusen

Abstract

Suppose a bidder must decide whether and when to incur the cost of estimating his own private value in an auction. This can explain why a bidder might increase his bid ceiling in the course of an auction, and why a bidder would like to know the private values of other bidders. It also can explain sniping—flurries of bids at the end of auctions with deadlines— as the result of other bidders trying to avoid stimulating the uninformed bidder to examine his value.

Indiana University Foundation Professor, Department of Business Economics and Public Policy, Kelley School of Business, Indiana University, 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. JEL Code: D44.

I thank Michael Fishman, David Hirshleifer, Michihiro Kandori, Roger Lagunoff, Dan Levin, James Peck, Michael Rothkopf, Jeffrey Stake, Lixin Ye, and participants at seminars at Ohio State and Tokyo Universities for their comments, and Ariel Kemper for research assistance. 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 Ebay 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 Ebay 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 using information from the other players' bids to estimate the maximum amount he would be willing to pay for the objection being auctioned. Rather, his problem is just to figure out how to minimize the amount he pays while still winning the item if the winning bid is less than his personal value. In practice, however, private-

¹From a story related to me by my law school colleague, Jeffrey Stake. I have taken artistic license with details.

value auctions can be even more difficult for the bidder than common-value auctions. This is for an “engineering reason” absent from our 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 as 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.

Fortunately, the bidder in an ascending private-value auction usually does not 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 why bidding in second-price private-value auctions is so easy. Under tight competition, however, even figuring out whether one’s value is the highest in the crowd is hard. “Value discovery”, as I shall call it, becomes an important part of the problem and explain a number of odd features we observe in real-world auctions:

1. Why bidders would like to know how much other bidders are going to bid.
2. How other bidders can benefit when an uninformed bidder learns his value more precisely.
3. How improved buyer information on the value of the object can hurt the seller.
4. Why bidders update the bid ceilings they submit during the course of an auction such as those on Ebay and Amazon that uses proxy bidding.
5. Why bidders use “pre-emptive bids”, bidding early in auctions rather than later.
6. Why bidders use “sniping”— the practice of submitting bids at the last minute.

Value discovery is not the only way to explain such things as updating of bids, pre-emptive bids, and sniping. 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 (Michael Fishman [1988], David Hirshleifer & Ivan Png [1989]). A variety of papers try to explain sniping, based on common values (Patrick Bajari & Ali Hortacsu [2000]), uncertainty over whether late bids will be registered by the auctioneer (Alvin Roth and Axel Ockenfels [2001]), and irrationality (Deepak Malhotra & Keith Murnighan [2000]). And the cost of returning to bid in an auction that takes place over several days has been shown by Octavian Carare and Michael Rothkopf (2001) 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 idea of value discovery has some similarity to the idea of entry fees in auctions. In the model of Dan Levin and James Smith (1994) 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, discovery will occur during the auction, raising the opportunity for strategic behavior by other bidders. 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 a different style, Nicola Persico (2000) has studied the efficiency of incentives to acquire information, and shows that with independent private values the Vickrey mechanism is efficient.²

The closest model to the present one is that of Oliver Compte and Philippe Jehiel (2000), who compare ascending and second-price sealed-bid private-value auctions when there are n bidders and one of them can acquire information on his value during the course of the auction. The difference from the current paper is that they focus on the effect of the number of bidders

²Information acquisition in common value auctions is a separate topic, since it has a public good aspect. See Donald Hausch and Lode Li (1993) for such a model.

and compare expected welfare and revenue from the two types of auction rules, and they do not consider the “hard” and “soft” ending rules that we will examine below. They also do not require any time elapse between the decision to discover the value and the completion of the discovery, an assumption which can be very important in ascending auctions.

In another stream of research 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. In their model, bids do stimulate discontinuous behavior, as will the model of the present paper, but each player knows his own valuation, so what the pre-empting player is trying to do is to show the other player that the pre-empting player’s valuation is high so as to make him give up on winning.³

I will proceed by laying out a model of value discovery (Section 2), which I will analyze under three sets of auction rules, starting with a second-price sealed-bid auction in order to most simply show the tradeoffs involved (Section 3). I will then lay out the equilibria of two kinds of open auctions, the Amazon and Ebay auctions (Section 4) and conclude (Section 5).

2. The Model

The two players in an auction are both risk-avertal and have private values which are statistically independent and distributed over the same support $[0, \bar{u}]$, on which the densities are strictly positive. Bidder 1 (the “uninformed” bidder) has value u_1 distributed according to the atomless and differentiable density $f(u_1)$. Bidder 2 (the “informed” bidder) has value u_2 distributed

³Another stream of research, quite different from this one but perhaps easily confused is where the asymmetric information in a private-value auction concerns the number of bidders, which one bidder may or may not be able to ascertain at a cost. See Kenneth Burdett and Kenneth Judd (1983), R. Preston McAfee and John Macmillan (1987), and Richard Engelbrecht-Wiggins (2000) for such models.

according to the atomless and differentiable density $g(u_2)$. Bidder 1 knows neither u_1 nor u_2 . At any time he may, unobserved by Bidder 2, pay c and learn u_1 after additional time δ has passed.⁴ Bidder 2 knows u_2 , but not u_1 .

The model is designed to address a situation in which one bidder is uncertain about (a) his value and (b) whether some other bidder has a higher value. This bidder's decision whether to discover his value is the driving force of the model.

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.

3. The Second-Price Sealed-Bid Auction

In the sealed-bid second-price auction, each player submits one “bid ceiling”, without knowing what the other has done. Whoever submits the highest bid ceiling wins the auction, but pays the bid ceiling submitted by the other player, or zero if the other player chose not to participate.⁵

3a. Bidding Strategies

*Equilibrium.*⁶ Bidder 1 pays to discover his value and submits a bid ceiling of

⁴The assumption that Bidder 2 does not observe Bidder 1's payment of c is made for simplicity, not to drive results. If it were not made, then we would have to specify out-of-equilibrium beliefs for Bidder 2 about Bidder 1's type if he were to deviate and irrationally choose to pay c at the wrong time.

⁵I do not allow the seller to post a reserve price, even though he might profit by using one under certain specifications of $f(u_1)$ and $g(u_2)$. This allows the analysis to keep its focus on buyer behavior without imposing tighter conditions on the value distributions.

⁶I use “equilibrium” to mean perfect Bayesian equilibrium. Second-price sealed-bid auctions do have equilibria other than the one described here, but they are perverse ones in which players use weakly dominated strategies. Consider a private-value auction with two bidders, A and B. When it is common knowledge that Bidder A will win a second-

u_1 if c is sufficiently low. Otherwise, he submits a bid ceiling of Eu_1 . Bidder 2 submits a bid ceiling of u_2 .

Bidder 1's Bidding Strategy. First, suppose Bidder 1 has paid c and discovered his value, u_1 . Once Bidder 1 knows u_1 , if Bidder 2 submits a bid ceiling of p with probability $m(p)$, Bidder 1's expected payoff is

$$\pi_1(u_1) = \int_0^b (u_1 - p)m(p)dp. \quad (1)$$

Maximizing by choice of b yields $(u_1 - b)m(b) = 0$, so $b^* = u_1$. Bidder 1 should bid his value, u_1 .

Second, suppose Bidder 1 has not discovered his value. His payoff if he bids b and Bidder 2 bids p with probability $m(p)$ is

$$\pi_1 = \int_0^{\bar{u}} \left(\int_0^b (u_1 - p)m(p)dp \right) f(u_1)du_1. \quad (2)$$

Maximizing by choice of b yields

$$\int_0^{\bar{u}} (u_1 - b)m(b)f(u_1)du_1 = 0, \quad (3)$$

so

$$\int_0^{\bar{u}} bm(b)f(u_1)du_1 = \int_0^{\bar{u}} u_1m(b)f(u_1)du_1, \quad (4)$$

and $b^* = \int_0^{\bar{u}} u_1f(u_1)du_1$, the expected value of u_1 . Bidder 1 should bid his expected value.

Bidder 2's Bidding Strategy. Bidder 2's payoff if he submits a bid ceiling of b and Bidder 1 submits a bid ceiling of p with probability $m(p)$ is

$$\pi_2 = \int_0^b (u_2 - p)m(p)dp, \quad (5)$$

price auction, he is willing to bid higher than his value, v_A , and Bidder B is willing to bid less than v_B . That will also be true in the Ebay and Amazon auctions to be considered here. Thus, if $v_A = 3$ and $v_B = 10$, one Nash equilibrium in the second-price auction is for Bidder A to bid 100 and Bidder B to bid 0. These strategies are weakly dominated, however. Bidder A would do no worse with a bid of 3 and would do better if Bidder B bid 90. Bidder B would do no worse with a bid of 10, and would do better if Bidder A bid 9. I will thus ignore these perverse equilibria.

because Bidder 2 wins the value u_2 and pays the price p if his bid of b exceeds Bidder 1's bid of p , and otherwise his payoff is zero. Maximizing his payoff by choice of b yields

$$(u_2 - b)m(b) = 0, \tag{6}$$

so $b^* = u_2$. Bidder 2 should bid his value.

3b. Bidder 1's Decision of Whether to Pay to Discover His Value

Let us now turn to Bidder 1's decision of whether to pay to discover his value. Denote his expected payoff at the start of the game from the strategy of paying to discover his value by π_1^d and his expected payoff from the strategy of not paying to discover his value by π_1^{nd} .

First, suppose Bidder 1 chooses the strategy of paying c to discover his value, u_1 . His payoff will be either $-c$, if $u_1 < u_2$ and he loses the auction, or $-c + u_1 - u_2$, if he wins the auction. Bidder 1 does not know Bidder 2's value, u_2 , but we can nonetheless calculate Bidder 1's expected payoff conditional upon u_2 . For given Bidder 2 value u_2 , Bidder 1's expected payoff before he actually learns u_1 is, integrating over all values of u_1 ,

$$\pi_1^d(u_2) = -c + \int_0^{u_2} (0)f(u_1)du_1 + \int_{u_2}^{\bar{u}} (u_1 - u_2)f(u_1)du_1. \quad (7)$$

The second term, which equals zero, arises because Bidder 1 loses the auction if $u_1 < u_2$. The third term, which is positive because u_1 ranges between values of u_2 and z in the integral, arises because he wins if $u_1 > u_2$.

For high enough values of u_2 , $\pi_1^d(u_2) < \pi_1^{nd}$: Bidder 1's payoff from paying c to learn his value is less than his payoff from not paying c (and in fact $\pi_1^d(u_2) < 0$ for large enough u_2). This is because the third term of $\pi_1^d(u_2)$ shrinks to zero as u_2 increases: for large u_2 , there is little chance Bidder 1 will want to overbid Bidder 2. Thus we obtain Proposition 1.

Proposition 1: *Bidder 1 benefits from knowing Bidder 2's value; if c is low enough, Bidder 1's expected payoff is higher if he learns u_2 before deciding whether to learn u_1 .*

Bidder 1 benefits from learning u_2 in two ways. First, if $\pi_1^d > \pi_1^{nd}$, Bidder 1 would switch from always paying to discover his value to paying only if $\pi_1^d(u_2) > \pi_1^{nd}$. Second, if $\pi_1^d < \pi_1^{nd}$, Bidder 1 would switch from never discovering his value to discovering it if $\pi_1^d(u_2) > \pi_1^{nd}$. If c is sufficiently high, however, then knowing u_2 is useless to Bidder 1, because for every u_2 , $\pi_1^d(u_2) < \pi_1^{nd}$. In that situation, Bidder 1 would never pay to discover u_1 even if he knew u_2 .

In the standard private-value second-price auction model, knowing the values, or even the bids of other bidders is unhelpful, because the knowledge would not affect one's strategy.⁷ Its usefulness here is that it does affect his decision about learning his own value more precisely.

Let us now return to deriving Bidder 1's expected payoff for the strategy of paying to discover his value when he does not know u_2 . Integrating over all possible values of u_2 yields an overall expected payoff of

$$\pi_1^d = -c + \int_0^{\bar{u}} \left(\int_{u_2}^{\bar{u}} (u_1 - u_2) f(u_1) du_1 \right) g(u_2) du_2. \quad (8)$$

For comparison with his payoff when he does not acquire information, it will be useful to divide the integral in this payoff into two parts, depending on whether u_2 is less than Eu_1 or greater, and represent them by A_1 and A_2 . Both A_1 and A_2 are positive, since both represent situations in which Bidder 1 wins the auction at a price less than his value.

$$\begin{aligned} \pi_1^d &= -c + \int_0^{Eu_1} \left(\int_{u_2}^{\bar{u}} (u_1 - u_2) f(u_1) du_1 \right) g(u_2) du_2 + \int_{Eu_1}^{\bar{u}} \left(\int_{u_2}^{\bar{u}} (u_1 - u_2) f(u_1) du_1 \right) g(u_2) du_2 \\ &= -c + A_1 + A_2. \end{aligned} \quad (9)$$

Now let us find Bidder 1's payoff if he does not learn u_1 . He will bid Eu_1 . If $Eu_1 < u_2$, he will lose the auction and his payoff will be zero: $\pi_1^{nd}(u_2) = 0$. If $Eu_1 > u_2$, he will win, and his payoff will be $u_1 - u_2$. Integrating over the possible values of u_1 , it is useful to divide the integral into two parts as follows:

$$\pi_1^{nd}(u_2) = \int_0^{u_2} (u_1 - u_2) f(u_1) du_1 + \int_{u_2}^{\bar{u}} (u_1 - u_2) f(u_1) du_1. \quad (10)$$

The first integral is negative, because Bidder 1 is winning the auction and paying u_2 when $u_1 < u_2$. The second integral is positive, because $u_1 > u_2$.

⁷This is not true in all models. Jeitschko (1998) models a sequential first-price sealed-bid auction in which bidders have either high or low private values, so that ties are likely. The equilibrium is in mixed strategies, and knowing the values of one's rivals is helpful in deciding one's bid.

Integrating $\pi_1^{nd}(u_2)$ over u_2 yields the expected payoff

$$\begin{aligned}
\pi_1^{nd} &= \int_0^{Eu_1} \left(\int_0^{u_2} (u_1 - u_2) f(u_1) du_1 + \int_{u_2}^{\bar{u}} (u_1 - u_2) f(u_1) du_1 \right) g(u_2) du_2 + \int_{Eu_1}^{\bar{u}} (0) g(u_2) du_2 \\
&= \int_0^{Eu_1} \left(\int_0^{u_2} (u_1 - u_2) f(u_1) du_1 \right) g(u_2) du_2 + \int_0^{Eu_1} \left(\int_{u_2}^{\bar{u}} (u_1 - u_2) f(u_1) du_1 \right) g(u_2) du_2 \\
&= -A_3 + A_1,
\end{aligned} \tag{11}$$

where A_1 is the same integral A_1 as in expression (9). The term $-A_3$ is negative because it represents the outcomes in which Bidder 1 has won and paid $u_2 > u_1$.

These results illustrate the benefits of information. If c is too large, then $\pi_1^d < 0$, and value discovery is inferior to non-discovery. If c is small enough, however, it is worthwhile to pay it to discover u_1 . This is true because Bidder 1's expected payoff from discovering his value is $\pi_1^d = -c + A_1 + A_2$, while if he does not discover his value it is $\pi_1^{nd} - A_3 + A_1$. If c is small enough, discovery has the higher expected payoff because it includes the benefit A_2 and avoids the cost $-A_3$.

All three terms have intuitive meanings. Term A_1 , present in both π_1^d and π_1^{nd} , is the payoff from winning profitably when u_2 takes the low values between zero and Eu_1 . The strategy of discovery adds term A_2 , the payoff from winning profitably when u_2 takes values higher than Eu_1 (values which still have some probability of being less than Bidder 1's value, u_1). The strategy of non-discovery incurs the extra cost, $-A_3$, the payoff from winning unprofitably when u_2 takes low values between zero and Eu_1 . Thus, discovery has the two benefits of winning profitably more often, and of never winning unprofitably.

Thus, Bidder 1 will pay to discover his value if c is low enough, and otherwise not pay, as asserted in the proposed equilibrium.

3c. Value Discovery and the Payoffs of Bidder 2 and the Seller

So far we discussed the value of information to Bidder 1. We know that Bidder 1 benefits from knowing his value. How does Bidder 1's value discovery affect the payoffs of the Bidder 2 and the seller? This is addressed by Propositions 2 and 3, which apply not only to the second-price sealed-bid auction but to all three auction rules.

Proposition 2: *The expected payoff of the informed bidder (Bidder 2) is higher if the uninformed bidder (Bidder 1) knows u_1 at the start of the auction.*

Proof: First, suppose $u_2 < Eu_1$. Bidder 2's payoff is zero without value discovery, because Bidder 1 will bid Eu_1 and win. Bidder 2's payoff is positive with value discovery, because there is probability $F(u_2)$ that u_1 will be less than u_2 and Bidder 1 will win. Thus, if $u_2 < Eu_1$, Bidder 2 benefits from Bidder 1 knowing u_1 .

Second, suppose $u_2 \geq Eu_1$. Without value discovery, Bidder 2's payoff is

$$u_2 - Eu_1, \quad (12)$$

which can be rewritten as

$$\int_0^{\bar{u}} u_2 f(u_1) du_1 - \int_0^{\bar{u}} u_1 f(u_1) du_1. \quad (13)$$

With value discovery, Bidder 2's payoff is

$$\int_0^{u_2} (u_2 - u_1) f(u_1) du_1 + \int_{u_2}^{\bar{u}} (0) f(u_1) du_1 = \int_0^{u_2} u_2 f(u_1) du_1 - \int_0^{u_2} u_1 f(u_1) du_1. \quad (14)$$

We need to show that (13), Bidder 2's payoff when Bidder 1 does not discover his value, is less than (14), Bidder 2's payoff when Bidder 1 does discover his value, so we need to show that

$$\int_0^{\bar{u}} u_2 f(u_1) du_1 - \int_0^{\bar{u}} u_1 f(u_1) du_1 < \int_0^{u_2} u_2 f(u_1) du_1 - \int_0^{u_2} u_1 f(u_1) du_1, \quad (15)$$

which is equivalent to

$$\int_{u_2}^{\bar{u}} u_2 f(u_1) du_1 < \int_{u_2}^{\bar{u}} u_1 f(u_1) du_1, \quad (16)$$

which is true because in the right-hand-side integral u_1 is taking values that are u_2 or greater.

Hence, Bidder 2 benefits from Bidder 1's value discovery at the start of the auction, which was to be shown.⁸

When Bidder 1 knows his value, there is a gain in total surplus, because the auction becomes efficient; whichever bidder has the higher value wins the auction. Bidder 1 would benefit from knowing u_1 at zero cost, as shown in Section 3b, and we have just seen from Proposition 2 that Bidder 2 benefits from Bidder 1 knowing u_1 . Since there is an efficiency gain, it is also possible that the seller gains from improved information about the value of his good to Bidder 1. Proposition 3 says that this turns out not to be the case.

Proposition 3: *The seller can prefer that a bidder (Bidder 1 here) not know his value precisely at the start of the auction.*

Proof: Take a given u_2 . First, suppose $u_2 < Eu_1$. The winning price would be u_2 . Value discovery will either keep the winning price at u_2 (if $u_1 \geq u_2$), or reduce it to below u_2 (if $u_1 < u_2$). Thus, the expected winning price is higher if Bidder 1 does not know u_1 .

Second, suppose $u_2 \geq Eu_1$. The winning price would be Eu_1 if Bidder 1 does not know u_1 . Value discovery will change the winning price to $\text{Min}(u_1, u_2)$. The winning price is u_1 if $u_1 < u_2$ and u_2 if $u_1 > u_2$, so its expected value is

$$\int_0^{u_2} u_1 f(u_1) du_1 + \int_{u_2}^{\bar{u}} u_2 f(u_1) du_1. \quad (17)$$

This is less than Eu_1 if

$$\int_0^{u_2} u_1 f(u_1) du_1 + \int_{u_2}^{\bar{u}} u_2 f(u_1) du_1 < \int_0^{\bar{u}} u_1 f(u_1) du_1, \quad (18)$$

⁸I include the caveat "at the start of the auction" because in the Amazon and Ebay auctions value discovery may occur later, with different results for the informed bidder and the seller.

which is true if

$$\int_{u_2}^{\bar{u}} u_2 f(u_1) du_1 < \int_{u_2}^{\bar{u}} u_1 f(u_1) du_1, \quad (19)$$

which is true since in the second integral u_1 ranges from u_2 up to \bar{u} whereas in the first integral u_2 is a constant.

Thus, the winning price falls with value discovery for any given value of u_2 , which was to be shown.

Proposition 3 is in striking contrast to the conventional wisdom about auctions, which is that the seller should do everything possible to improve information about the object's value (see Milgrom and Weber [1982], or, for a recent elaboration showing how a seller might strategically manipulate information, Kaplan and Zamir [2000]). That conventional wisdom is, of course, true in its proper context. In a common value auction, the better the information, the less acute is the winner's curse. Moreover, if buyers are risk averse, whether in a private value auction or a common value one, then reducing the uncertainty they have over the object's value increases their willingness to bid. And as recent work by Lixin Ye (2001a, 2001b) has shown, if bidders must bear an entry cost prior to participating in the auction, then the seller should take into account the effect of improved information on how many bidders will enter. Indeed, the seller may wish to use an "indicative bidding" round prior to the main auction, constructed so that bidders will reveal something of their values to each other.

In the present model, however, none of these considerations apply. Instead, the expected price falls when the bidder knows his private value better. The average bid ceiling submitted by Bidder 1 is Eu_1 whether he knows his value or not. If he does not know his value, however, his bid ceiling is always Eu_1 with certainty, whereas if he does know his value, his bid ceiling has greater variation and the expectation of the minimum of his bid ceiling and Bidder 2's bid ceiling becomes smaller.

The result here is also a point of difference between the present paper and the model of Compte and Jehiel (2000), who find that if the number

of bidders is large enough and the discovery cost is small enough, then the seller's revenue is higher in an ascending auction (in which information is revealed in the course of the auction) than in a second-price sealed-bid auction. An interesting feature of their model is that an uninformed bidder should not pay to discover his value so long as at least two other bidders are active in the bidding—which means that he will remain in the auction even after the price has exceeded his expected value, because he may decide later (instantaneously) to pay to discover his value. If there are many other bidders, the uninformed bidder is unlikely to have the highest value, and so runs little risk of winning at a low price and later discovering that his value is even lower. Rather, if he discovers his value, it is likely to be after the winning bid has exceeded his expected value, and so the effect of value discovery is to raise his bid rather than to lower it. This points to an important difference in the effects of value discovery on seller incentives when there are many instead of few bidders, and when the uninformed bidder is unlikely instead of likely to be the winner.

4. Open-Cry Auctions

We will now look at open-cry auctions, in which the bidders can make deductions about each other's values in the course of the auction by observing the bidding.

The Amazon auction uses “proxy bidding”. Each player submits a “bid ceiling”. The “current winner” is the player with the highest bid ceiling, who if nobody revised bid ceilings would win the auction and pay the second-highest bid ceiling, the “current winning bid” (which is set at 0 at the start of the auction). The auctioneer publicly posts the current winning bid, but not the bid ceilings. At any time, a player can increase his bid ceiling, but he cannot reduce it. The auction ends either at time T , or N minutes after the last bid ceiling revision, whichever is later, with $N > \delta$, so there is still time for value discovery. Thus, the auction has a “soft deadline.” The current winner becomes the winner and pays the current winning bid. If there is a tie, then the current winners have equal probabilities of becoming the winner.

The Ebay auction also uses proxy bidding. It is the same as the Amazon auction except that the auction ends at time T regardless of when the last bid ceiling revision took place. Thus, the auction has a “hard deadline.”

4a. The Amazon Auction

Because the Amazon auction has a soft deadline, Bidder 1 always has time to discover his value before he must respond to Bidder 2’s bids.

Equilibrium. Bidder 1 submits a bid ceiling of either Eu_1 , if c is high enough, or $\bar{p} > 0$ otherwise. If he has submitted a bid ceiling of \bar{p} and the current winning bid rises to \bar{p} , he pays c to discover u_1 and then increases his bid ceiling to u_1 if $u_1 > \bar{p}$. Bidder 2 submits a bid ceiling of u_2 .

Explanation. In the Amazon auction, as in the second-price sealed bid auction, the winner is the bidder with the highest bid ceiling and the price equals the second-highest bid ceiling. Thus, the only difference is that the bidders can extract information by observing the current winning bid.

It will be convenient to discuss Bidder 2’s strategy first. For the same reasons as in the second-price sealed-bid auction, Bidder 2 will want to have submitted a bid ceiling of u_2 by the end of the auction at time T : bidding less than u_2 , he would lose when he could profitably win while not reducing the price he paid; bidding more than u_2 he might win at a price greater than his value. The only question is whether Bidder 2 could benefit from choosing the timing of his bid so as to affect Bidder 1’s value discovery decision. In the Amazon auction, Bidder 2 cannot so benefit. Whether Bidder 2 submits a bid ceiling of \bar{p} early in the auction or late, Bidder 1 has time for value discovery, because when Bidder 2 submits such a bid ceiling the current winning bid increases and the soft deadline means that a time interval of at least $N > \delta$ remains in the auction. Thus, Bidder 2 submits a bid ceiling of u_2 at any time before T .

For the same reasons as in the second-price sealed-bid auction, Bidder 1’s optimal bid ceiling at the end of the auction is u_1 if he knows u_1 and Eu_1

otherwise. If Bidder 1 chooses the policy of submitting a bid ceiling of Eu_1 and never paying to discover u_1 , his expected payoff is identical to equation (10), its value in the second-price sealed-bid auction. Just as in that auction, he wins at the price of u_2 if $u_2 < Eu_1$ and loses if $u_2 > Eu_1$.

If Bidder 1 chooses the policy of submitting a bid ceiling of \bar{p} , his expected payoff is different. Let us again construct a $\pi_1^d(u_2)$ function, now for the strategy of discovering u_1 if the current winning bid reaches \bar{p} . If $u_2 < \bar{p}$, Bidder 1's expected payoff is

$$\pi_1^d(u_2) = \int_0^{\bar{u}} (u_1 - u_2) f(u_1) du_1. \quad (20)$$

If $u_2 \geq \bar{p}$, Bidder 1's expected payoff is

$$\pi_1^d(u_2) = -c + \int_0^{u_2} (0) f(u_1) du_1 + \int_{u_2}^{\bar{u}} (u_1 - u_2) f(u_1) du_1. \quad (21)$$

Integrating over the possible values of u_2 yields an overall expected payoff for Bidder 1 of

$$\pi_1^d = \int_0^{\bar{p}} \left(\int_0^{\bar{u}} (u_1 - u_2) f(u_1) du_1 \right) g(u_2) du_2 + \int_{\bar{p}}^{\bar{u}} \left(-c + \int_{u_2}^{\bar{u}} (u_1 - u_2) f(u_1) du_1 \right) g(u_2) du_2. \quad (22)$$

If $\bar{p} = 0$, Bidder 1 pays to discover u_1 immediately as the auction starts. The first term of payoff (22) equals zero, leaving the payoff identical to equation (8), the same as in the second-price sealed-bid auction for the policy of discovering u_1 . Thus, our results from the second-price sealed-bid auction tell us immediately that $\pi_1^d > \pi_1^{nd}$ if c is low enough but not otherwise. All that remains to be shown is that he will choose $\bar{p} > 0$, a strategy in which he discovers u_1 only if the current winning bid rises above zero.

Lemma: *Bidder 1 will set his initial bid ceiling to be strictly positive, delaying value discovery: $\bar{p} > 0$.*

Proof: Differentiating Bidder 1's payoff with respect to \bar{p} (which we can do

if the densities are atomless) yields

$$\frac{d\pi_1^d}{d\bar{p}} = \left(\int_0^{\bar{u}} (u_1 - \bar{p})f(u_1)du_1 \right) g(\bar{p}) - cg(\bar{p}) - \left(\int_{\bar{p}}^{\bar{u}} (u_1 - \bar{p})f(u_1)du_1 \right) g(\bar{p}) \quad (23)$$

If $\bar{p} = 0$, the first and third terms of this derivative cancel out. The second term is positive, however, given our assumption that the value density is everywhere positive. Thus the payoff derivative is positive at $\bar{p} = 0$, in which case the optimal value of \bar{p} is positive, which was to be proved.

There is an intuitive explanation for this result. The advantage of increasing \bar{p} is that possibly the bidder will win at a price $p < \bar{p}$ and not have to pay the discovery cost c . The disadvantage of increasing \bar{p} is that possibly the bidder will win at a price $p < \bar{p}$ such that p exceeds his value: $p = u_2 > u_1$, something which value discovery would have prevented. The size of this disadvantage depends on the likelihood that Bidder 1's value is below \bar{p} , which is $\int_0^{\bar{p}} f(u_1)du_1$. If $\bar{p} = 0$, this disadvantage vanishes; there is no risk that Bidder 1's value will be below \bar{p} . Thus, Bidder 1 should increase his initial bid ceiling until the marginal gain from avoiding the discovery cost equals the marginal loss from winning when his value is below the price he pays.

In the equilibrium of the Amazon auction, Bidder 1 raises his bid ceiling if it happens that $u_2 \geq \bar{p}$. This yields us Proposition 4.

Proposition 4: *The phenomenon of bidders increasing their bid ceilings during the auction can be a necessary part of equilibrium.*

Without value discovery, a bidder has no incentive to wait to submit his true value as his bid ceiling, rather than submit it at the start of the auction. In the present model, that would be true if the bidder either knew u_1 at the start or if he could not spend c to discover u_1 . Here, however, a bidder can find it strictly superior 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 that

bidding is competitive enough.⁹

Proposition 4 could arise in another way, not present in this model, but true to its spirit: 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. 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.”¹⁰

4b. The Ebay Auction

The Ebay auction is identical to the Amazon auction except that it has a hard deadline: bid ceilings cannot be updated after time T regardless of when the current winning bid last increased. This means that Bidder 1 cannot discover his value in response to bids observed after time $T - \delta$.

Equilibrium. If c is low enough, Bidder 1 submits a bid ceiling of \bar{b} before time $T - \delta$, and then pays to discover his value if the current winning bid reaches \bar{b} . If the current winning bid does not reach \bar{b} , he raises his bid ceiling to Eu_1 . If c is higher, Bidder 1 submits a bid ceiling of Eu_1 and never

⁹I explore this point further in Rasmusen (2002), where I use it to explain the phenomenon of bidders apparently getting “carried away” in auctions.

¹⁰“As eBay Rivals Emerge, Some Tips on Bidding And Selling on the Web,” *Wall Street Journal*, Bart Ziegler, September 23, 1999, p. B1.

discovers his value. Bidder 2 submits a bid ceiling of \bar{b} before time $T - \delta$, and raises his bid ceiling to u_2 after time $T - \delta$.

Explanation. For the same reasons as in the first two auction rules, in the Ebay auction Bidder 1's bid ceiling by the end of the auction will be u_1 if he knows it and Eu_1 otherwise, and Bidder 2's bid ceiling will be u_2 . The only questions are whether Bidder 1 wants to submit a lower bid ceiling in order to learn something about Bidder 2's value, and whether Bidder 2 is willing to act so as to reveal his value.

The Ebay auction has introduced a new complication. Now Bidder 2 can wait to bid until after $T - \delta$ and prevent value discovery. Whether Bidder 2 wants to do this depends on whether he wants to provoke Bidder 1 to discover u_1 . If Bidder 2 has value u_2 and provokes value discovery by submitting a bid ceiling of \bar{b} or more, his expected payoff is made up of three parts, depending on whether he wins at a price of \bar{b} , wins at a price of u_1 , or loses the auction.

$$\pi_2 = \int_0^{\bar{b}} (u_2 - \bar{b})f(u_1)du_1 + \int_{\bar{b}}^{u_2} (u_2 - u_1)f(u_1)du_1 + \int_{u_2}^{\bar{u}} (0)f(u_1)du_1. \quad (24)$$

This payoff is declining in \bar{b} , since its derivative is

$$\frac{d\pi_2}{d\bar{b}} = (u_2 - \bar{b})f(\bar{b}) - \int_0^{\bar{b}} f(u_1)du_1 + (u_2 - \bar{b})f(\bar{b}). \quad (25)$$

Intuitively, if \bar{b} is greater, then Bidder 2 benefits less from Bidder 1 learning u_1 because Bidder 1 cannot reduce his bid ceiling below \bar{b} after learning that $u < \bar{b}$.

If $\bar{b} = 0$, expression (24) becomes identical to Bidder 2's payoff when Bidder 1 discovers u_1 at the start of the auction, the value found earlier in equation (14). We can deduce from Proposition 2 that there is some positive value of \bar{b} low enough that Bidder 2 will benefit from value discovery, since if $\bar{b} = 0$, Proposition 2 tells us that Bidder 2 benefits from value discovery. Let us call that critical value k .

In choosing a discovery threshold, \bar{b} , Bidder 1 is constrained to set $\bar{b} \leq k$, because otherwise Bidder 2 would want to avoid provoking value discovery

and would delay submitting u_2 as his bid ceiling until after $T - \delta$. Thus, \bar{b} will be less than or equal to the \bar{p} of the Amazon auction. If c is too high, then Bidder 1 will prefer not to discover his value and to simply bid Eu_1 at the start of the auction.

In the equilibrium of the Ebay auction, it is important for Bidder 2 to submit a bid ceiling of at least \bar{b} before time $(T - \delta)$ if he wishes to provoke value discovery. This yields us Proposition 5.

Proposition 5: *A bidder may purposely bid early, so as to stimulate value discovery.*

Proposition 5 stands in contrast to the standard private-value auction model, in which the timing of bids is unimportant. If there were no value discovery in the model— or if value discovery could occur at any time, as in the Amazon auction— then it would make no difference whether Bidder 2 bid early or late. In the Ebay auction, however, it is important that Bidder 1 submit the bid ceiling of \bar{b} before time $T - \delta$, because he wishes to stimulate value discovery. Bidder 1 has submitted a relatively low bid ceiling, and Bidder 2 would like him to discover u_1 rather than just raise his bid ceiling to Eu_1 . Bidder 2 thus makes what looks like a pre-emptive bid, but not for the usual reason of deterring entry into the auction when there is an entry cost, but for what is almost the opposite reason: to stimulate another bidder to buy costly information and perhaps leave the auction as a result. Bidder 2 wishes to alert Bidder 1 that there will be tight competition to buy the object being sold, so that Bidder 1 will think carefully before he advances his bid ceiling to Eu_1 .

Proposition 5 also shows that sniping is not always the optimal strategy— indeed, in this version of the model, it is never an optimal strategy. It has been noted that if there is a common value component to the value, then a bidder ought to bid as late as possible, in order to prevent his rivals from learning about their values from his bid (see, e.g., p. 369 of Wilcox [2000]). Here in this private-value setting, however, the effect of bidding early can be to reduce competition rather than intensify it.

4c. The Equilibrium if Bidder 1 is Naive

Let us call Bidder 1 a “naive bidder” if he places zero probability on Bidder 2 being in the auction and at the start of the auction submits his expected value Eu_1 , as his bid ceiling. In such a situation, of course, Bidder 1 expects to win at a price of zero (since he expects no other bids to be made), so he is indifferent as to his bid ceiling. To submit his expected value, however, is to follow the advice of auction theory, which says that submitting one’s private value is a dominant strategy, and is to follow the advice of the Ebay website instructions, which say:

“For example, if the current bid on an item is \$5 and you are willing to pay up to \$10, you would enter \$10 as your maximum bid. Your bid would be shown on the item page as \$5, but if another bidder places a bid for \$6, then eBay will place a higher bid on your behalf. The bid would be just above the other member’s bid. This would continue until either you win the auction at or below \$10 or the bidding exceeds the \$10 you were willing to pay. eBay will notify you via email if you are outbid and you can return to place another bid if you like. Your maximum bid is never disclosed to other bidders or to the seller.” — EBay Tutorials, “Place Your Bid”

http://pages.ebay.com/education/tutorial/course1/bidding_3.html (May 25, 2002)

Ronald Wilcox (2000) investigated the pattern of bidding in Ebay auctions for consumer items and found that less experienced bidders submitted their bids earlier than more experienced bidders; for example, 1.2 percent of the least experienced bidders bid during the last minute, whereas 8.2 percent of the most experienced bidders did. This effect was present for both private-value and common-value goods, though most pronounced for common-value goods.¹¹

¹¹Wilcox also tried to approach the question of whether less experienced bidders increased their bids during the course of the auction. This is difficult because the public data does not say whether a bidder has increased his bid. It does, however, have the total number

Alvin Roth and Axel Ockenfels (2001, 2002) have also investigated Internet auctions. They find that EBay auctions have more sniping than Amazon auctions— 20 percent of all last Ebay bids are submitted in the last hour, but only 7 percent of Amazon bids, and that in the last five minutes 9 and 16 percent (for computers vs. antiques) of Ebay bidders submit their bids but only 1 percent of Amazon bidders. The difference is even clearer with experienced bidders: in the Ebay auction they submit bids later than the inexperienced bidders, while on Amazon they submit earlier. They also surveyed bidders who bid in the last minute. Most of the some 70 bidders who replied said they bid late consciously and to keep prices down. Some said they were influenced in their values by the bids of other people, but 88 percent of late bidders said that they had a clear idea of their value at the start of the auction. A few (less than 10 percent) seem to have been confused about the auction rule and thought they had to bid late to win.

In the second-price sealed-bid auction, Bidder 2 would still submit a bid ceiling of u_2 if Bidder 1 is naive and bids Eu_1 . As a result, for low c , the expected price would be higher than when Bidder 1 was not naive. seller. Bidder 1's naivete has prevented him from paying c to discover his value, to the detriment of both bidders and to the advantage of the seller.

In the Amazon auction, if Bidder 1 starts by submitting a bid ceiling of Eu_1 , Bidder 2 would also submit a bid ceiling of u_2 . The current winning bid would then equal $\text{Min}(Eu_1, u_2)$, and Bidder 1 would realize that the probability-zero event of Bidder 2 being present at the auction had occurred. If the current winning bid were less than Eu_1 , Bidder 1 would realize that $Eu_1 > u_2$ and not pay to discover his value. If the current winning bid were Eu_1 , Bidder 1 would realize that if he does not discover his value, his payoff will equal zero because he will lose the auction, but if he does discover his value, his expected payoff will be

$$-c + \int_{Eu_1}^{\bar{u}} \left(\int_0^{u_2} (0)f(u_1)du_1 + \int_{u_2}^{\bar{u}} (u_1 - u_2)f(u_1)du_1 \right) g(u_2)du_2. \quad (26)$$

of bids and bidders. He found no significant difference in number of bids per bidder in auctions for private-value goods where average experience was higher, but that the number of bids per bidder was higher for common-value goods.

Whether this payoff is positive or negative depends on the value of c , so Bidder 1 will discover his value if c is small enough and not otherwise. Bidder 2 has no reason to delay submitting a bid ceiling of u_2 because under the Amazon auction rule, Bidder 2 has time $N > \delta$ after the current winning bid rises to Eu_1 in which to discover his value. As in the second-price sealed-bid auction, the expected price is higher when Bidder 1 is naive, to the detriment of Bidders 1 and 2 and to the benefit of the seller.

Bidder 2 has no reason to delay submitting his bid ceiling of u_2 because under the Amazon auction rule, Bidder 2 has time $N > \delta$ after the current winning bid rises to Eu_1 in which to discover his value. He would like to alert Bidder 1 to his presence before the auction starts, in order to prevent Bidder 1 from starting with a bid ceiling as high as Eu_1 , but that is not possible. As in the second-price sealed-bid auction, the expected price is higher when Bidder 1 is naive, to the detriment of Bidders 1 and 2 and to the benefit of the seller.

In the Ebay auction, if the naive Bidder 1 would submit a bid ceiling of Eu_1 at the start of the auction, Bidder 2's best response is slightly more complicated. If $u_2 \leq Eu_1$, Bidder 2 would submit a bid ceiling of u_2 at any time. If $u_2 > Eu_1$, then if expression (26) is positive, Bidder 2 would submit a bid ceiling of u_2 in the time interval $[T - \delta, T]$; if it is negative, he would submit a bid ceiling of u_2 at any time.

The difference between the Amazon and Ebay auctions is that in the Ebay auction, Bidder 2 can avoid stimulating Bidder 1's value discovery. The difference between the Ebay auction in Section 4b and the Ebay auction with the naive bidder is that the naive bidder has bid Eu_1 , which is above the critical level k below which Bidder 2 wants to stimulate value discovery. That $Eu_1 > k$ is simple to see. If Bidder 1 discovers u_1 after submitting a bid ceiling of Eu_1 , either he discovers that $u_1 < Eu_1$ and does not change his bid ceiling, or he discovers that $u_1 > Eu_1$ and increases his bid ceiling. When Bidder 1 increases his bid ceiling, Bidder 2 is hurt, either because Bidder 2 still wins but at the increased price of $u_1 > Eu_1$ or because now Bidder 2 loses because $u_1 > u_2$. Thus, Bidder 2 wants to avoid stimulating value

discovery, which the submission of a bid ceiling of u_2 before $T - \delta$ would do if $u_2 > Eu_1$ and expression (26) were positive. When Bidder 1 is naive in the Ebay auction, Bidder 2 will to delay submitting his bid ceiling if $u_2 > Eu_1$, which yields us Proposition 6.

Proposition 6: *Sniping can occur in equilibrium. A bidder may purposely delay submitting a bid ceiling higher than the current winning bid until near the auction deadline.*

Sniping would not occur in a standard private-value auction model, in which the timing of bids does not matter, because other bidders will not find it useful to learn one's value. When a bidder is uncertain about his private value, however, as in the present model, he does benefit from learning other bidders' values (Proposition 1) and this can be to their benefit (Proposition 5) or detriment (Proposition 6).

Other Explanations for Sniping

The simplest alternative explanation for sniping is 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.

This explanation— though, of course, phrased less technically— can be found with others at the website, “advanced Auction Management,” http://www.tblightning.com/ebay/auction_management.htm (Jan 27, 2002), which lists other practical reasons for bidding late. One such reason is that a late bidder does not have to commit early to buy that item only to later find something he would rather buy instead. Closely related is that the late bidder knows whether he has successfully bought the item quickly, so he can move on to another source if he loses. Neither of these reasons is a compelling one for bidding in the last moments of the auction, but the site also mentions one

that is: “shilling” sellers. Sellers in Internet auctions are strictly forbidden to bid on their own items, although they may use a pre-set reserve price if they wish. Otherwise, the seller could see how much the highest-valuing bidder was willing to pay and turn the auction into a bargaining game to increase the price. Nonetheless, shilling is hard to catch, because sellers can use pseudonymous email addresses to bid. If such strategizing by the seller takes time, however, it can be evaded by bidding close to the deadline. This is a plausible risk, though it seems unlikely that seller fraud is so rampant as to explain the level of sniping that we observe.

Alvin Roth and Axel Ockenfels (2001) 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 Ebay, he finds that other bidders will outbid him and the resulting bidding war will leave even the winner with a low payoff.¹²

Compte and Jehiel (2000) do not consider sniping in their paper, but their model could be adapted to explain it in a certain context. They construct a model of an ascending private-value auction in which value discovery is costly but instantaneous, and they find that if there are sufficiently many informed bidders, expected seller revenue is higher than in a second-price sealed-bid auction. As explained above, revenue rises because the uninformed bidder is unlikely to have the highest value, and will wait until late in the auction to decide whether to discover his value. Thus, as in the naive bidder

¹²This 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.

model in the present paper, value discovery will raise his bid, not reduce it. Suppose, instead, that value discovery requires time and that the auction is an Ebay auction, with a hard deadline. The informed bidders may then have incentive to delay until the last few minutes of the auction so that the uninformed bidder cannot use their bids to decide whether to increase his bid above his expected value. This explanation applies to when there are many informed bidders and the uninformed bidder is unlikely to have the highest bid; the value discovery model in the current paper applies when there are few informed bidders and the uninformed bidder has a bigger chance of having the highest bid.

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 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.¹³

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 Ebay 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.

¹³The idea that finding out one's own value for an object explains odd behavior also shows up in bargaining. In Rasmusen (2001), 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.

References

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

Burdett, Kenneth & Kenneth Judd (1983) “Equilibrium Price Dispersion,” *Econometrica*, 51: 981-996 (July 1983).

McAfee, R. Preston & John Macmillan (1987) “Auctions with a Stochastic Number of Bidders,” *Journal of Economic Theory*, 43: 1-19.

Carare, Octavian & Michael Rothkopf (2001) “Slow Dutch Auctions,” RUTCOR Research Report RRR42-2001, Rutgers University.

Compte, Olivier & Philippe Jehiel (2000) “On the Virtues of the Ascending Price Auction: New Insights in the Private Value Setting,” working paper, CERAS-ENPC, CNRS, Paris (December 2000).

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).

Engelbrecht-Wiggans, Richard (2001) “The Effect of Entry and Information Costs on Oral versus Sealed-Bid Auctions,” *Economics Letters*, 70, 2:195-202 (February 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).

Hausch, Donald & Li, Lode (1993) “A Common Value Auction Model with Endogenous Entry and Information Acquisition,” [*Economic Theory*, 3(2): 315-334 (April 1993)].

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 (2002) "Bertrand Competition Under Uncertainty," *Journal of Industrial Economics*, 50: 11-21 (March 2002).

Jeitschko, Thomas (1998) "Learning in Sequential Auctions," *Southern Economic Journal*, 65: 98-112. (July 1998).

Kaplan, Todd & Shmuel Zamir (2000) "The Strategic Use of Seller Information in Private-Value Auctions," Hebrew University, Center For Rationality Working Paper No. 221 (July 2000).

Larson, Kate & Tuomas Sandholm (2002) "Costly Valuation Computation in Auctions," Carnegie Mellon University, Dept. of Computer Science, undated.

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.

Milgrom, Paul and Robert Weber (1982) "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50:1089-1122.

Persico, Nicola (2000) "Information Acquisition in Auctions," *Econometrica*, 68: 135-148 (January 2000).

Rasmusen, Eric (2001) "Explaining Incomplete Contracts as the Result of Contract-Reading Costs," *Advances in Economic Analysis and Policy*, Vol. 1: No. 1, Article 2,
<http://www.bepress.com/bejeap/advances/vol1/iss1/art2>.

Rasmusen, Eric (2002) "Getting Carried Away in Auctions as Imperfect Value Discovery,"

[Http://Php.indiana.edu/~erasmuse/papers/carried.pdf](http://Php.indiana.edu/~erasmuse/papers/carried.pdf).

Roth, Alvin E. & Axel Ockenfels (2001) “Last Minute Bidding And The Rules For Ending Second-Price Auctions: Theory And Evidence From A Natural Experiment On The Internet,” working paper, Harvard University Dept. of Economics, June 7, 2001.

Roth, Alvin E. & Axel Ockenfels (2002) “Last Minute Bidding And The Rules For Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet,” forthcoming, *American Economic Review*.

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

Ye, Lixin (2001a) “Optimal Auctions with Endogenous Entry,” working paper, Ohio State Dept. of Economics (June 1, 2001).

Ye, Lixin (2001b) “Indicative Bidding,” working paper, Ohio State Dept. of Economics (July 1, 2001).