

Getting Carried Away in Auctions as Imperfect Value Discovery

June 6, 2004

Eric Rasmusen

Abstract

Bidders in auctions must decide whether and when to incur the cost of estimating the most they are willing to pay. This can explain why people seem to get carried away, bidding higher than they had planned before the auction and then finding they had paid more than the object was worth to them. Even when such behavior is rational, *ex ante*, it may be perceived as irrational if one ignores other situations in which people revise their bid ceilings upwards and are happy when that enables them to win the auction.

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, www.rasmusen.org. www.rasmusen.org/papers/carried-rasmusen.pdf.

I thank Michael Baye, Rene Kirkegaard, Kerry Krutilla, Yacheng Sun, and participants in seminars at Indiana University's Kelley School of Business, DePaul University, the University of Tokyo's economics department, and the 2004 International Industrial Organization Conference for their comments, and Harvard Law School's Olin Center and the University of Tokyo's Center for International Research on the Japanese Economy for their hospitality.

1. Introduction

Do people get carried away at auctions? Certainly they do in the sense of bidding more than they had intended when they arrived at the auction. What must be determined is why that happens. Most simply, other people bid more than expected, so a bidder ends up paying more than he thought he would. But what we usually think of as getting carried away is that during the auction a bidder raises the maximum amount he is willing to bid.

Even this often has a simple explanation: that the auction is for an object with a common value component, so that during the auction our bidder learns from the high bids of other people that the object is worth more than he had thought earlier. That, too, is not what we think of as getting carried away. Rather, we think of a private value auction— of some good that the bidder intends to keep for personal consumption rather than for production or resale— in which the bidder ends up bidding more than his pre-auction estimate of the value of the object.

The standard advice to bidders is to avoid getting carried away. One website puts it like this:

Never go beyond a predetermined limit when bidding. Base this limit on the information you have gathered. Avoid becoming obsessed with an item. Doing so will lead you to bid more than the property (or merchandise) is worth. If you are bidding on a tax sale property, you might bring a certified cheque for the maximum amount you intend to bid. This should ensure that you do not get carried away with the bidding process. If you are the successful bidder and the property is sold for less than the amount on your cheque, the clerk/treasurer will issue a refund for the difference.

Avoid catching auction fever. This happens when bidders get carried away with the process; they will bid on anything and everything that is being auctioned and often will end up being the owner of things they did not even want and paying far too much for these items. The opposite of auction fever is auction paralysis. This occurs when the bidder is paralysed with fear and thus is unable to make a bid. Apparently such a state is often due to a fear of overpaying. If you don't overcome it you will never get started.

Often, if you fail to do your homework, you will not have the confidence to bid. (“Tax Sale Properties/Auction Guidelines,” (http://www.taxsaleproperties.com/abt_7.html))

It is quite plausible that people make such mistakes (see Malhotra & Murnighan [2000] for a persuasive example of irrational bidding, or the survey evidence of confusion about auction rules on page 14 of Roth & Ockenfels [2002]), and even more plausible that in auctions, as in ordinary purchases, people end up spending extravagantly on current consumption to their later regret. I wish to propose another explanation, however, for bidders who end up paying more than their pre-auction maxima: that the bidder rationally revises his estimate of the value of the object upwards during the course of the auction, so at the moment of purchase he actually does value it at more than the price he pays. The model will still be one with purely private values, since our bidder will not be deriving any information about his own value from the other players’ bids. The difference from the standard private-value auction will be that it is costly for him to discover his own private value, so he defers doing so until the middle of the auction. At that point he might revise it upwards—“auction fever”—or he might revise it downwards—“auction paralysis”.

The model that will be used is similar to the models of Compte & Jehiel (2000) and Rasmusen (2003a): a bidder in a private-value ascending auction will be uncertain about his value and will be able to pay a fixed amount to improve his information. Unlike in those models, here his information will still be imperfect after value discovery, and we will focus on his decision to update his bid rather than on the auction’s payoffs or welfare implications under different auction rules. Note, too, that the situation is quite unlike that in models such as Persico (2000) and the articles cited there which examine the incentives to gather information on the value of objects only before, not during, an auction.

2. The Model

There are two possible buyers in an auction for an object, both risk-

neutral and with private values which are statistically independent.

Bidder 1's value is v_1 , which has three components: $v_1 = \mu + u + \epsilon$. Bidder 1 does not know the sizes of u , ϵ , or v_1 . He does know μ , and he knows that u and ϵ are independently distributed according to symmetric densities $f(u)$ and $g(\epsilon)$ with mean zero and supports such that $\text{Min}(\mu + u + \epsilon) \geq 0$, so that $v_1 \geq 0$. As a result, Bidder 1's initial expectation of v_1 equals μ . If he wishes, at any time he can pay c and learn the value of u immediately. He cannot discover the other component, ϵ , however, until after the auction.

Bidder 2's value, v_2 , is α with probability θ and β with probability $(1 - \theta)$, where α , θ , and β are common knowledge; and $\theta \in (0, 1)$, $\alpha \in (0, \mu)$, and $\beta > \mu$. We do not need any assumption on the expected value of v_2 relative to v_1 . Bidder 2 knows the value of v_2 but not the value of v_1 .

The auction is open and ascending. The price starts at zero and rises continuously until one player drops out, at which point the other player wins the object and pays that price. A player's bidding strategy is a choice of a price at which to drop out (a "bid ceiling") in an open- exit auction such as this one. This setup avoids the technical untidiness created when the winning player must bid a positive increment higher than the next-highest bid in order to win.

Discussion of the Assumptions

Our purpose is to model a situation in which a bidder is uncertain about (a) his value and (b) whether there exists any other player whose value is higher. The model's focus is on his decision on whether to incur the cost of learning more about his value.

We assume the low value for Bidder 2, α , to be less than μ so that if $v_2 = \alpha$ Bidder 1 will win the auction even if he just bids up to μ . The value α is assumed to exceed $\mu - \underline{u}$ so that Bidder 1's ex post payoff might be negative if he wins at a price of α .

We assume the possible high value for Bidder 2, β , exceeds μ so that

if $v_2 = \beta$ Bidder 1 will lose the auction if he just bids up to μ . Note that however high β may be, under our assumptions there is still some chance that $(\mu + u)$ will be high enough that Bidder 1 will win the auction.

A different way to model this situation would be to assume general differentiable distributions for v_2 as well as for u and ϵ . That model will be used in Section 4. The assumption of a two-point distribution used here, however, will allow for some interesting comparative statics, and will be heuristically useful.

In this game it is especially important to think of the probabilities of each of Bidder 2's types, $v_2 = \alpha$ and $v_2 = \beta$, as being the subjective probabilities of the uninformed player, which are not necessarily the true population magnitudes. The variable θ represents the strength of Bidder 1's belief that Bidder 2's value is low. Note that θ can be arbitrarily close to one and the results of the model still hold. The model is most interesting for high values of θ , which represent situations in which Bidder 1 is surprised to find that he faces tight competition from Bidder 2.

Rasmusen (2003a) also models a bidder who begins a private value auction unsure of his own value but who can pay c to acquire information. The important differences between that model and this one are that here if the uninformed bidder pays c then (a) he acquires the value information immediately, not after a time lag, and (b) he only acquires better information about his value, not perfect information. The absence of a time lag means that the informed bidder has no incentive to strategically delay bidding, the "sniping" phenomenon at the heart of my other article. The imperfection of the information means that even if the uninformed bidder makes optimal decisions ex ante, ex post he may regret having made them.

3. The Equilibrium

Each bidder must decide on a bid ceiling. Bidder 1 must also decide at what bid level, if any, to pay c to discover u , after which he may wish to revise his bid ceiling.

The optimal bidding strategies are straightforward. A player should choose a bid ceiling equal to the expected value of the object being auctioned. If he bids any less, he could lose even though the winning price was less than his expected value. If he bids any more, he could win at a price greater than his expected value.

Thus, if Bidder 1 does not acquire any information about his value, his best strategy is to bid up to μ , the expected value of the object to him. If he does discover u , his optimal strategy is to bid up to $(\mu + u)$, his updated estimate of v_1 . Bidder 2's optimal strategy is to choose a bid ceiling of v_2 . Note that there is no benefit to Bidder 2 in changing his bid ceiling in order to affect the timing of Bidder 1's value discovery; value discovery is instantaneous, so timing is unimportant in this model, unlike in Rasmusen (2003a), where value discovery cannot take place late in the auction.

Bidder 1 has three value discovery strategies that might be optimal in equilibrium: early discovery, late discovery, and no discovery. The early discovery strategy is to pay to discover u when the bid level reaches some value $b^* \in [0, \alpha)$, most simply at the start of the auction, so $b^* = 0$. The late discovery strategy is to pay to discover u if the bid level reaches some level $b^* \geq \alpha$ and Bidder 2 has failed to drop out, most simply if the bidding reaches Bidder 1's initial bid ceiling, so $b^* = \mu$. The strategy of no discovery is to refuse to pay to discover u regardless of what happens.

Bidder 1's expected payoff if he chooses never to pay to discover his value is

$$\pi_1(\textit{no discovery}) = \theta(\mu - \alpha) + (1 - \theta)(0), \quad (1)$$

because with probability θ he will win the auction at a price of α and with probability $(1 - \theta)$ he will lose the auction.

Bidder 1's expected payoff from late discovery— paying c and discovering v_1 if and only if Bidder 2 does not drop out at the price of α — is made up of the expected value of winning at price α and the expected value of

winning at price β when $\mu + u > \beta$.

$$\pi_1(\text{late discovery}) = \theta(\mu - \alpha) + [1 - \theta] \left[-c + \int_{\beta - \mu}^{\infty} [(\mu + u) - \beta] f(u) du \right] \quad (2)$$

Bidder 1's payoff from early discovery is made up of the cost c plus the expected value of winning at price α when $\mu + u > \alpha$ and $v_2 = \alpha$, plus the expected value of winning at price β when $\mu + u > \beta$ and $v_2 = \beta$.

$$\begin{aligned} \pi_1(\text{early discovery}) = & \theta \left[-c + \int_{\alpha - \mu}^{\infty} [(\mu + u) - \alpha] f(u) du \right] + \\ & [1 - \theta] \left[-c + \int_{\beta - \mu}^{\infty} [(\mu + u) - \beta] f(u) du \right] \end{aligned} \quad (3)$$

We can now make more precise what happens in equilibrium. If c is low enough, early discovery is best: Bidder 1 pays c at the start of the game to avoid the negative payoff from winning even at the low price of α when v_1 happens to be low. For moderate levels of c , late discovery is best: he waits until he observes Bidder 2 bidding more than α and then to pay c to discover his value. If c is high enough, no discovery is best; he sticks with his initial bid ceiling and never pays c . The optimal value discovery strategies are made precise in Proposition 1, which uses the following shorthand notation:

$$\begin{aligned} A_1 & \equiv \int_{\alpha - \mu}^{\infty} [(\mu + u) - \alpha] f(u) du \\ A_2 & \equiv \int_{\beta - \mu}^{\infty} [(\mu + u) - \beta] f(u) du \end{aligned} \quad (4)$$

Proposition 1: *All three value discovery strategies can be optimal, depending on the value of the discovery cost c :*

$$\begin{aligned} \text{Early discovery} & \quad \text{if } c \leq A_1 - (\mu - \alpha) \\ \text{Late discovery} & \quad \text{if } c \in [A_1 - (\mu - \alpha), A_2] \\ \text{No discovery} & \quad \text{if } c \geq A_2 \end{aligned} \quad (5)$$

If Bidder 2's possible high value is farther from μ than his possible low value, then the middle range is empty and late discovery is not optimal: If $\beta - \mu > \mu - \alpha$ then late discovery will not occur.

Proof: The payoffs in equations (3), (2), and (1) can be written as:

$$\begin{aligned}
(a) \quad \pi_1(\text{early discovery}) &= \theta(-c + A_1) + (1 - \theta)(-c + A_2) \\
(b) \quad \pi_1(\text{late discovery}) &= \theta(\mu - \alpha) + (1 - \theta)(-c + A_2) \quad (6) \\
(c) \quad \pi_1(\text{no discovery}) &= \theta(\mu - \alpha)
\end{aligned}$$

The early discovery payoff in equation (6a) is greater than the late discovery payoff in equation (6b) whenever $c < A_1 - (\mu - \alpha)$. The no discovery payoff in equation (6c) is greater than the late discovery payoff in equation (6b) whenever $c > A_2$.

If $A_1 - (\mu - \alpha) < A_2$, then the middle range of c values is not empty and any of the three strategies might be optimal. That inequality can be rewritten as

$$\int_{\alpha - \mu}^{\infty} (\mu + u - \alpha) f(u) du - (\mu - \alpha) < \int_{\beta - \mu}^{\infty} (\mu + u - \beta) f(u) du \quad (7)$$

or

$$\begin{aligned}
0 &< -(\mu - \alpha)[F(\infty) - F(\alpha - \mu)] - \int_{\alpha - \mu}^{\beta - \mu} u f(u) du + (\mu - \alpha) + (\mu - \beta)[F(\infty) - F(\beta - \mu)] \\
0 &< (\mu - \alpha)F(\alpha - \mu) - \int_{\alpha - \mu}^{\beta - \mu} u f(u) du + (\mu - \beta)[1 - F(\beta - \mu)].
\end{aligned} \quad (8)$$

Suppose $\mu - \alpha > \beta - \mu$. Then the middle term of the right side of inequality (8), $-\int_{\alpha - \mu}^{\beta - \mu} u f(u) du$, which is the expected value of u conditional on being between $\mu - \alpha$ and $\beta - \mu$, is positive. Is the positive first term's magnitude greater than the negative third term's? Yes, for the following reason. Since the distribution F is symmetric, if $\mu - \alpha = \beta - \mu$ (contrary to our assumptions) then

$$(\mu - \alpha)F(\alpha - \mu) + (\mu - \beta)[1 - F(\beta - \mu)] \quad (9)$$

equals zero. Suppose now that we increase $\mu - \alpha$, so that it exceeds $\beta - \mu$. Since $d/dz zF(z) = F(z) + zf(z) > 0$, this increases the first, positive, term of expression (9) while leaving the second, negative, term unaffected. Thus, expression (9), when brought into accordance with the assumption that $\mu - \alpha > \beta - \mu$, is positive, expressions (8) and (7) are positive, and the middle range of c in Proposition 1 is not empty; otherwise, it is, and late discovery is not optimal. ■

Inequality (7) is true if $\mu - \alpha > \beta - \mu$; that is, if the low Bidder 2 value, α , is further from Bidder 1's expected value, μ , than is the high Bidder 2 value, β . If this were not true, then as the discovery cost c increased, Bidder 1 would simply jump from early discovery to no discovery. Bidder 1 would choose no discovery if c became too big to justify paying it to avoid overpaying α when $\mu + u > \alpha$. But in that case, when the bid rose to α and Bidder 2 was still in the auction, Bidder 1 would find *a fortiori* that c was too big to justify paying it to gain the chance of winning the auction when $\mu + u > \beta$.

Having established the equilibrium, we can now see how Bidder 1's behavior is affected by changes in parameters other than c .

Proposition 2: *Bidder 1's willingness to pay to improve his estimate of falls with the toughness of competition but is unaffected by the probability of tough competition: the level of c which makes "No Discovery" optimal is falling in β and unchanged in θ .*

Proof: Bidder 1's willingness to pay to improve his estimate is captured by the bounds in Proposition 1. In particular, he will follow the policy of no discovery if and only if $c \geq A_2$, which is to say, if

$$c \geq \int_{\beta-\mu}^{\infty} [(\mu + u) - \beta]f(u)du \quad (10)$$

The derivative of A_2 with respect to β is then

$$\frac{dA_2}{d\beta} = -(\mu + [\beta - \mu] - \beta)f(\beta - \mu) - \int_{\beta-\mu}^{\infty} f(u)du, \quad (11)$$

which is negative. The derivative of A_2 with respect to θ is zero.

Thus, increases in β expand the parameter range for no discovery but increases in θ leave it unchanged. ■

The intuition behind the first part of Proposition 2 is that as β increases and becomes further from μ , it becomes less likely that the expected value after discovery, $(\mu + u)$, will be greater than β and Bidder 1 will want to increase his bid and win the auction. Thus, giving up becomes more attractive, unless the cost c of discovering u is low. It is interesting to see what happens near $\beta = \mu$ (though the assumptions of the model rule out $\beta \leq \mu$). For $\beta \leq \mu$, Bidder 1 would not pay even a tiny c to discover u , because Bidder 2 would have already dropped out and it would be too late for Bidder 1 to change his behavior. If, however, β is just slightly above μ , then the value of information about u is very large because with a probability of almost .5, discovery of u will lead Bidder 1 to change his behavior.

The key to the second part of Proposition 2 is that the decision between late and no discovery is deferred until new information arrives that renders irrelevant θ , the probability that $v_2 = \beta$. Bidder 1 need not decide about paying c until he sees that Bidder 2's value must be high— at which point the ex ante probability it is high is moot.

Interpretation as Getting Carried Away

This model provides an interpretation for “getting carried away” in an auction. Suppose we see a bidder winning an auction at a price higher than the most he entered the auction being willing to pay, and that he later regrets having won at that high price—what I will call an “unhappy victory.” At the start of the auction, μ was the most Bidder 1 intended to bid. The auction begins, and the bidding rises to μ . Now, however, he reconsiders, and raises his bid ceiling to $(\mu + u)$. This new ceiling is greater than β , the most Bidder 2 will pay. Bidder 1 thus wins the auction, at price β . After the auction is over, however, he discovers ϵ and finds that $\mu + u + \epsilon < \beta$. He says to himself: “I got carried away and bid too much. I wish I'd stuck with my

original ceiling of μ .”

This, of course, is only one possible scenario. It is worth exploring the conditions under which unhappy victories occur. In the story above, Bidder 1 had an unhappy victory. With equal likelihood, after the auction is over he would have discovered $\epsilon > 0$, so his consumer surplus would have been even higher than he had expected— a sort of “extra-happy victory”.

It is worth comparing the probability of unhappy victories with and without value discovery. Even if he does not discover u and increase his bid ceiling, Bidder 1 will still sometimes overpay. Suppose we are in the middle range of costs, so Bidder 1 is following the policy of late discovery. If Bidder 2’s value turns out to be low ($v_2 = \alpha$), Bidder 1 will win at a price of α . This is less than the expected value of v_1 , which is μ , but it might be more than the true value of v_1 , which is $\mu + u + \epsilon$, giving rise to an unhappy victory. Since u and ϵ have symmetric distributions, however, unhappy victories will occur with probability less than 50%. Indeed, if α is low, unhappy victories may be very rare. Victories without value discovery— without “getting carried away”— will occur only if the competition from Bidder 2 is weak, so the winning price is low and Bidder 1 will come away with a good chance of sizeable consumer surplus.

On the other hand, if Bidder 2’s value turns out to be high ($v_2 = \beta$), Bidder 1 will pay to discover u , and if u is high enough he will raise his bid ceiling high enough to win at a price of β . This is less than the expected value of v_1 , which is $\mu + u$, but it might be more than the true value of v_1 , which is $\mu + u + \epsilon$, giving rise to an unhappy victory.

What is the probability of this unhappy victory? For the borderline case of $\mu + u = \beta$, an unhappy victory has probability 50%. For higher values of u , the probability of an unhappy victory falls. But victories following value discovery— those that happen because the bidder “gets carried away”— occur only if the competition from Bidder 2 is strong. The winning price is thus high, and Bidder 1’s chance of coming away with positive consumer surplus may be very little higher than his chance of coming away with negative consumer surplus.

Thus, situations in which Bidder 1 pays to discover his value and then increases his bid ceiling and wins the auction are worse situations for him than when he does not pay to discover his value but wins the auction anyway. This is not because the strategy of late discovery is suboptimal, though— it is not. Rather, it is because actually carrying through and discovering his value under that strategy only occurs after bad news— the news that he is facing tough bidding competition. The strategy of late discovery is analogous to a person’s strategy of using chemotherapy if he is diagnosed with cancer. Under that strategy, chemotherapy will come to be associated with pain and death, but that does not lessen its usefulness in making the best of a bad situation.

4. A Model with Continuous Densities

In the model above, Bidder 2 had two possible values. This brings sharply into relief the late discovery strategy, in which Bidder 1 delays discovering his value in the hope that nobody else will have a high value. Another possible case, not more general, but equally interesting, is when Bidder 2 has a continuous distribution for his value. We will model that in a way similar to Rasmusen (2003a), adapted to the auction rules and instantaneous value discovery assumed in the present paper. This will show that the phenomenon of a bidder increasing his reservation bid in the course of an auction is robust, and, indeed, in the continuous density model not only will value discovery be optimal if c is not too large, but the optimality of late discovery for moderate levels of c will not require any condition analogous to that stated in Proposition 1.

As in Section 2, let there be two possible bidders, both risk-neutral, with private values which are statistically independent.

Our assumption about Bidder 1 will remain the same. Bidder 1’s value is v_1 , which has three components: $v_1 = \mu + u + \epsilon$. Bidder 1 does not know the sizes of u , ϵ , or v_1 . He does know μ , and he knows that that u and ϵ are independently distributed according to symmetric densities $f(u)$ and $g(\epsilon)$ with mean zero and supports such that $Min(\mu + u + \epsilon) \geq 0$, so that $v_1 \geq 0$.

As a result, Bidder 1's initial expectation of v_1 equals μ . If he wishes, at any time he can pay c and learn the value of u immediately. He cannot discover the other component, ϵ , however, until after the auction.

Unlike in Section 2, we will now assume that Bidder 2's value, v_2 , is distributed according to an atomless and differentiable density $h(v_2)$ on $[\alpha, \beta]$, where $0 < \alpha < \mu$ and $\beta > \mu$ and where $h(v_2) > 0$ for all v_2 on that interval. Bidder 2 does not know v_1 , but he does know v_2 . All parameters are common knowledge.

As in Section 3, Bidder 2's optimal strategy is to choose a bid ceiling equal to v_2 . Bidder 1's optimal bid ceiling is Ev_1 , which will be either μ or $(\mu + u)$, depending on whether he has learned u . Bidder 1 must also choose a "discovery level" p —a bid level at which Bidder 1 pays c to discover u , where possibly $p < \alpha$ (early discovery, because Bidder 1 will pay to discover his value before discovering anything about v_2) or $p > \mu$ (no discovery, because without discovery it never happens that the price rises above μ).

To analyze Bidder 1's payoff as a function of p , let us start by supposing (contrary to the assumptions) that Bidder 1 knows v_2 . Suppose also that $p \leq \mu$, so that there is positive probability that Bidder 1 pays c and discovers u .

If $v_2 < p$ then Bidder 1 wins the auction at price v_2 , for an expected payoff of $(\mu - v_2)$.

If $v_2 > p$ then he pays c to discover u . He loses the auction if $\mu + u < p$; otherwise, he wins. Overall, if $v_2 > p$ his expected payoff is

$$\pi_1(v_2|v_2 > p) = -c + \int_{u=-\infty}^{v_2-\mu} (0)f(u)du + \int_{u=v_2-\mu}^{\infty} (\mu + u - v_2)f(u)du. \quad (12)$$

Integrating over the possible values of v_2 yields an overall expected payoff for Bidder 1 of

$$\pi_1 = \int_{v_2=\alpha}^p (\mu - v_2)h(v_2)dv_2 + \int_{v_2=p}^{\beta} \left(-c + \int_{u=v_2-\mu}^{\infty} (\mu + u - v_2)f(u)du \right) h(v_2)dv_2 \quad (13)$$

If, on the other hand, $p > \mu$, then Bidder 1 is following the policy of no discovery, and his expected payoff is simply the first part of equation (13):

$$\pi_1(p > \mu) = \int_{v_2=\alpha}^{\mu} (\mu - v_2)h(v_2)dv_2. \quad (14)$$

Proposition 3: *In the model with continuous value densities, the optimal discovery level, p^* , rises with c , rising strictly if $p^* \in (\alpha, \mu)$. Bidder 1 will follow a policy of early discovery ($p^* \in [0, \alpha)$) if c is low enough, late discovery ($p^* \in [\alpha, \mu]$) for higher levels of c , and no discovery ($p^* \in (\mu, \infty]$) if c is sufficiently high.*

Proof: Differentiating equation (13) with respect to p yields

$$\begin{aligned} \frac{d\pi_1}{dp} &= (\mu - p)h(p) - \left(-c + \int_{u=p-\mu}^{\infty} (\mu + u - p)f(u)du\right) h(p) \\ &= \left[c + (\mu - p) - \int_{u=p-\mu}^{\infty} (\mu - p + u)f(u)du\right] h(p) \\ &= \left[c + - \int_{u=-\infty}^{\infty} (\mu - p + u)f(u)du - \int_{u=p-\mu}^{\infty} (\mu - p + u)f(u)du\right] h(p) \\ &= \left[c + \int_{u=-\infty}^{p-\mu} (\mu - p + u)f(u)du\right] h(p) \end{aligned} \quad (15)$$

If $h(p)$ is positive (which it is between α and μ) and c is small enough, then this derivative is negative. If c is small enough, $\frac{d\pi_1}{dp} < 0$ for $p \in [\alpha, \mu]$, and profit rises if p is reduced to below α – that is, to early discovery. If $p < \alpha$, then $h(p) = 0$, so further reductions are unimportant – early discovery can take the form of any p in the interval $[0, \alpha)$.

If c is greater, then $\frac{d\pi_1}{dp} > 0$ at $p = \alpha$, and the optimal p exceeds α .

Evaluated at the optimal level of p , the second derivative of profit is

$$\begin{aligned}
\frac{d^2\pi_1}{dp^2} &= \left((\mu - p + [p - \mu])f(u) + \int_{u=-\infty}^{p-\mu} (-1)f(u)du \right) h(p) + \left[c + \int_{u=-\infty}^{p-\mu} (\mu - p + u)f(u)du \right] h'(p) \\
&= 0 - \left(\int_{u=-\infty}^{p-\mu} f(u)du \right) h(p) + (0)h'(p), \\
&< 0,
\end{aligned} \tag{16}$$

where we use the fact that $\frac{d\pi_1}{dp} = 0$ at the optimum to obtain the term $(0)h'(p)$. Since it is also true that

$$\begin{aligned}
\frac{d^2\pi_1}{dpdc} &= (1)h(p) \\
&> 0,
\end{aligned} \tag{17}$$

the implicit function theorem tells us that $\frac{dp}{dc} > 0$ when $h(p) > 0$, i.e., the optimal discovery level rises continuously with the cost of discovery. This means that there exist levels of c such that the optimal discovery level lies within the interval (α, μ) , so late discovery is optimal. It also means that as c increases, eventually the optimal discovery level exceeds μ , so that no discovery becomes optimal. ■

If the discovery cost is low enough, early discovery is best, because the bidder averts the possibility that he might pay more than his value by winning even at the other bidder's lowest possible value. If the discovery cost is somewhat higher, it is not worth paying it to avoid that risk, and the bidder will choose late discovery. How late depends on the size of the discovery cost, and the optimal discovery level rises smoothly with the discovery cost, and if the discovery cost is too high, then no discovery becomes optimal.

Thus, we see that Section 3's conclusion that a bidder may decide to increase his bid ceiling in the course of an auction is robust to allowing rival values to take more than two possible levels. The comparative statics of Proposition 2 (on what happens when the probability θ of a low v_2 change)

are not easily adapted to the continuous model, but a new kind of comparative statics result possible only in the continuous model is derived in Proposition 4.

Proposition 4: *As the degree of uncertainty over his private value increases, Bidder 1 becomes more willing to pay to discover his value: if $p^* < \mu$, then p^* falls if we spread density $f(u)$ using a strict decrease in f on any interval $[r, s]$ and a strictly increase everywhere else, while leaving the mean of u unchanged at zero.*

Proof. Let us define X as a component of equation (15) (rearranged here slightly):

$$\begin{aligned} \frac{d\pi_1}{dp} &= (\mu - p)h(p) - \left(-c + \int_{u=p-\mu}^{\infty} (u - [\mu - p])f(u)du \right) h(p) \\ &= [\mu - p + c]h(p) - \left(\int_{u=p-\mu}^{\infty} Xf(u)du \right) h(p), \end{aligned} \tag{18}$$

where $x \equiv u - [\mu - p]$.

Changing f affects only the third term, which is always positive because it includes only values of u such that $(u - [\mu - p]) \geq 0$. Making f riskier using the conventional definition of risk from Rothschild & Stiglitz (1970) might leave X unchanged (as in Figure 1a, where the density changes in four regions, all to the right of $p - \mu$), or might increase it (as in Figure 1b, where the density falls between r and s but increases everywhere else).

Suppose, however, we spread out f using a strict decrease in f on any interval $[r, s]$ and a strictly increase everywhere else, while leaving the mean of u unchanged at zero. This is the continuous-distribution analog of the concept of “pointwise riskiness” that I explore at greater length in Rasmusen (2004). This flattens f because while the mean stays the same, the density strictly declines on the middle interval and strictly increases on each side of it. This forces an increase in f everywhere to the right of s . Since the spread leaves the unconditional mean of v_2 unchanged, it must increase the mean of v_2 conditional upon v_2 being above any specific value— and in particular, above $p - \mu$, so $\int_{u=p-\mu}^{\infty} Xf(u)du$ must increase. If $\int_{u=p-\mu}^{\infty} Xf(u)du$ increases,

then p must fall if we are at an interior solution and the derivative is to stay equal to zero. ■

Figure 1: Two Kinds of Increase in Risk

Proposition 4 is true because value discovery has option value, and option value increases with the amount of uncertainty. When the uncertainty is larger, there is a greater probability that value discovery will disclose that $\mu + u > \text{beta}$, even though the expected value of u is zero. Thus, even a risk-neutral Bidder 1 likes having more uncertainty.

Note that ϵ , the size of the remaining uncertainty over v_1 , is irrelevant to Bidder 1's decision. It could be that this uncertainty is far larger than that from u , but this makes no difference to the value to Bidder 1 of information about u . It would make a difference if Bidder 1 were risk averse, but we have assumed he is risk neutral, and a risk-neutral player only cares about variance to the extent that it affects option value. The variable $\bar{\epsilon}$ will enter our analysis in the next subsection.

5. Concluding Remarks

I have suggested an explanation for the phenomenon of bidders getting carried away in auctions. The man on the street would say that when a bidder has increased his bid ceiling from what he had decided before the auction, he has been overcome by emotion. He might also say that the bidder bid so high simply because he wanted to win, rather than because he wanted to own the object, but this variant too must rely on economic irrationality, since a rational bidder would factor his utility from winning into his original bid ceiling. I do not deny that there may be an emotional explanation; in fact, formalizing and testing such an explanation would be worthwhile. Here, however, I have proposed an alternative: in the course of the bidding, the bidder rethinks his private value, and with some probability his rethinking results in an upwards revision of the amount he is willing to pay. This revision is rational, and, indeed, it would be irrational for the bidder to incur too much cost in determining the maximum he is willing to pay before he knows whether that maximum will be a binding constraint.

Our paradigmatic example for the private-value auction is the open-cry ascending antique auction, but the idea of value discovery applies to any auction, and, indeed, the present model applies best when time pressure exists but is not so intense, so that a bidder does have a chance to reflect on his willingness to pay. One application with which many readers of this article may have experience is in house purchases. In buying a house, a person's first aim is to find a house for which his private value exceeds the likely price. Once he has found such a house, however, he may well find that other buyers also are interested in it, in which case an auction, usually informal but sometimes formal, begins. At that point, the buyer will think harder about his private value, and may revise it either up or down, but the buyer would have been foolish to undergo the emotional strain of such fine valuation if not forced to by tight competition.

The value discovery explanation for bid updates has three empirical implications that could help to test it. First, in the value discovery model the carried-away winner would regret having won less than half of the time—

less than half, because even though his revised value is still an overestimate, he usually will not have to pay the entire amount to win the auction. Second, a short “cooling off period” would presumably affect an emotional winner more than a value-discovering winner, although even in the value discovery model, the winner would, after thinking more, wish to return the object a significant fraction of the time.¹ Third, the value discovery model implies that if the value is more uncertain, the bidder will be more likely to increase his bid ceiling in the course of the auction, because the option value of value discovery is higher. An emotional explanation might or might not have this implication.

References

Asker, John (2000) “Bidding Up, Buying Out And Cooling-Off: An Examination Of Auctions With Withdrawal Rights” *Economic Theory*, 16: 585-611 (November 2000).

Bulow, Jeremy & Paul Klemperer (2002) “Prices and the Winner’s Curse,” *RAND Journal of Economics*, 33:1-21 (Spring 2002).

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).

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

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.

¹For an experimental study of cooling off periods which is aimed at seeing whether they benefit the seller by encouraging more aggressive bidding, see Asker (2000). See also the theoretical analysis of Rothkopf (1991).

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](http://www.bepress.com/bejeap/advances/vol1/iss1/art2).

Rasmusen, Eric (2003a) "Strategic Implications of Uncertainty Over One's Own Private Value in Auctions," January 26, 2003,
[Http://www.rasmusen.org/papers/auction.pdf](http://www.rasmusen.org/papers/auction.pdf).

Rasmusen, Eric (2004) "When Does Extra Risk Strictly Increase the Value of Options?" May 14, 2004,
[Http://www.rasmusen.org/papers/options.rasmusen.pdf](http://www.rasmusen.org/papers/options.rasmusen.pdf).

Roth, Alvin E. & Axel Ockenfels (2002) "Late-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon on the Internet," *American Economic Review*, 92(4): 1093-1103 (September 2002).

Rothkopf, Michael (1991) "Auctions With Withdrawable Winning Bids," *Marketing Science*, 10,1: 40-58 (Winter 1991)

Rothschild, Michael & Joseph Stiglitz (1970) "Increasing Risk I," *Journal of Economic Theory*, 2: 225-243 (September 1970).