

30
365
342 COPY 2

STX



BEBR

FACULTY WORKING
PAPER NO. 1342

THE LIBRARY OF THE
MAY 7 1987
UNIVERSITY OF ILLINOIS
URBANA

The Effect of Regret on Optimal Bidding in Auctions

Richard Engelbrecht-Wiggans

BEBR

FACULTY WORKING PAPER NO. 1342

College of Commerce and Business Administration

University of Illinois at Urbana-Champaign

March 1987


The Effect of Regret on Optimal Bidding in Auctions

Richard Engelbrecht-Wiggans, Associate Professor
Department of Business Administration

.

Abstract:

While most models of auctions and competitive bidding assume that each bidder's utility for an outcome depends only on his own profit, we allow the utility to also depend on any regret that a bidder suffers after the fact, for example over "money left on the table" in Federal offshore oil lease sales. Typically, for risk neutral bidders who, after the fact, know the winner's price for the object, a bidder's optimal bidding strategy will not depend on the relative weight given to profit versus regret. However, if losers do not learn the winner's price, then the bidders' reactions to regret hurts the bid-taker at equilibrium. Thus, the existing models' exclusion of regret from risk neutral bidders' utility functions affects the applicability of the resulting theory only under certain, now clearly delineated, conditions.



Digitized by the Internet Archive
in 2011 with funding from
University of Illinois Urbana-Champaign

Introduction:

Developments in the theory of auctions and competitive bidding derive from three sources--observations of the real world, experiments in the laboratory, and analyses of mathematical models. Each contributes in its own way. Each also has its own limitations. But, together they result in a richer theory than could result from any one source by itself.

Specifically, observing actual auctions carefully enough reveals real situations and phenomena to be studied, modelled, or explained. So, observations provide a necessary link to the real world. However, at best, we can only accurately observe what actually happens in actual situations as we perceive them. This affords little opportunity for repeated observations of what happens in a specific situations for defining the specific situation, or for modifying the specific situation.

On the other hand, laboratory experiments replace the real world with a model of the auction and its environment. This flexibility in defining the auction rules and environment allows for the observation of how the surrogates react to many different, not necessarily real world, situations. This flexibility also carries a cost--the responsibility to think about how the laboratory situation relates to the real world. In addition, experiments substitute surrogates for the real world bidders. This also carries a cost--the experiences, motivation, and expertise of the surrogates may differ from those of real bidders, and, therefore, so too might their reaction to any specific situation.

Finally, mathematical models go even one step further from the real world; they model not only the auction rules and the environment within which the auction occurs, but may model the bidders. If analytically tractable, such models suggest hypotheses to be tested experimentally, hypotheses that might otherwise not have been thought of. However, the increased flexibility in defining situations also carries an increased responsibility; by modelling the bidders in addition to the auction itself, the chances to misrepresent the real world increase.

In particular, the usefulness of a model depends on the extent to which it captures or represents something of practical interest. A model might give rise to certain phenomena observed in the real world, and thereby be useful in suggesting how and why such phenomena arise in the real world. Alternatively, a model might focus on one particular aspect of actual auctions, and thereby be useful in understanding how this aspect affects what happens in the real world and why. In either case, though, what practical insights can be gained from the model and its analysis depends on how sensitive the results are to exactly how the model approximates the real world.

This paper examines the sensitivity--or lack thereof--of the theory to a particular change in modelling the bidders themselves. By definition, each bidder bids as if he were maximizing his expected utility from the auction as he sees it; while this describes what bid a bidder will make, it does not prescribe, or even necessarily describe, how the bidder comes up with the bid. Of course, the bidder's utility depends on the outcome. The outcome, in turn, depends

on how everyone bids. Thus, a particular bidder's expected utility depends not only on his own bid, but also on how he believes others would bid in response to any possible view that they may have of the actual situation, as well as on our particular bidder's beliefs about the relative probability of others' views of the situation given his own view of the situation.

In some models, a bidder's expected utility maximizing bid remains the same regardless of how others bid; the expected utility may change, but not necessarily what bid maximizes the expected utility. However, in the absence of such "dominant bidding strategies," how a bidder bids depends on his perceptions about how others will bid. Still, if a bidder's perceptions about others changes very little from one auction to another, then the Nash equilibrium provides a possible characterization of everyone's behavior; an equilibrium exists if no one bidder could do better than to continue bidding as he has in the past so long as others continue to bid as they have in the past.

Not everyone accepts the Nash equilibrium as an appropriate model of bidder's behavior. Yet, despite any reservations, the Nash equilibrium remains a commonly used model. In fact, we know of no better model; in this sense at least, the Nash equilibrium is the natural model. Moreover, we suspect that most reservations with Nash equilibria actually stem more from an inappropriate choice of utility function or from an inappropriate model of how each bidder perceives others and the real world than from any flaw in the equilibrium concept itself. Therefore, we question the form of utility function typically used, or more precisely, we will examine the effects from modifying the usual utility function form.

Most existing models of auctions and competitive bidding assume one specific form of utility function. (For examples, see the work of Vickrey (1961), Myerson (1981), Milgrom and Weber (1982), or the survey of Engelbrecht-Wiggans (1980).) While the bidders--as modelled by the utility function--might be risk neutral or risk averse, the utility of any outcome depends only on the bidder's profit from that outcome. Although this captures what may be the most important component of bidders' utility functions, it ignores other potentially important components. We therefore look to the real world in asking what might we want to put into the utility function in addition to, or possibly in place of, profit.

In this spirit, reacting to comments from bidders on Federal offshore oil leases that they bid to maximize the quantity of mineral reserves won rather than to maximize expected profit--the profit on an oil lease being so affected by others' decisions once the lease has been won that the bidder might be unable to even define what is meant by "profit"--Engelbrecht-Wiggans (1987) defines a Principal-Agent model of competitive bidding. In the model, bidders act as agents for their respective oil firms, but bid to maximize how much they win gross of what they must pay for it. Of course, without any constraint on how bidders bid, such an objective would drive up the price of leases without limit. So, the oil firm, acting as the principal, places some constraint on how its agent may bid. For an appropriately set limit on bidder's expected expenditures, the bidder's bidding will be indistinguishable from how he would have bid were he maximizing expected profit. (Operationally, an oil firm may encourage its bidder to win as

many mineral reserves as possible subject to some longer term feedback on whether the bidder is on average spending too much or too little per sale.) On the other hand, any constraint on a bidder's total exposure in an auction typically results in distinctly different bidding. Thus, any theory for expected profit maximizing bidders applies equally well if bidders maximize expected gross winnings subject to an appropriate constraint on expected expenditures, but not if the constraint is on total exposure in a sale. This helps broaden and define the limits of the existing theory.

In the same spirit, the current paper reacts to the concern over "money left on the table"--the amount of money the winner could have saved himself in a first-price sealed-bid auction had he known, when he bid, what his nearest competitor would bid. This suggests letting the utility function depend on regret as well as on profit. Including the regret suffered by the loser who has a value for the object in excess of the price paid by the winner seems logical. Thus, keeping things as simple as possible, we consider defining a bidder's utility as a linear combination of profit and regret.

Roughly speaking, we shall show that if, at the end of the auction, a bidder learns the winner's price, then the relative weight given to profit versus regret in the utility function does not affect that bidder's optimum bidding strategy. Thus, all the theory for expected profit maximizing bidders applies equally well in many practical situations with bidders who are risk neutral but consider regret in addition to profit. However, a simple example illustrates the importance of bidders, after the fact, knowing the winner's price.

Bell's (1982) work in utility theory supports our choice of the utility function form. In particular, he suggests utility functions of the form $u(s,t) = v(s) + f(v(s)-v(t))$. Extending his definitions of s and t to our, more general, setting suggests indentifying s as a bidder's actual profit, and t as the profit that the bidder could have had from bidding different than he actually did had he known when he bid what he knows after the fact. Then for risk neutral bidders, the expression becomes $u(s,t) = \alpha s + f(\alpha s - \alpha t)$, where α is a constant between zero and one. Note that, since for almost all bids, the regret varies continuously with the bid, for our problem t must be no less than s . So defining $f(z) = (1-\alpha)z/\alpha$ for non-positive z , and $f(z) = 0$ for non-negative z gives $u(s,t) = \alpha s - (1-\alpha) \max \{0, (t-s)\}$; the utility is a linear combination of profit and regret, with the constant α parameterizing the weighting, just as we already suggested.

We close with a brief outline of what follows. The next section defines our notation, defines regret in terms of this notation, derives an expression for a bidder's ex-ante expected utility, and establishes conditions under which a bidder's optimal strategy is independent of the parameter α ; roughly speaking, if, after the fact, each bidder knows the winner's price, then the necessary condition will typically be satisfied, and so the optimal strategy will be independent of the parameter α . The following section provides two illustrative examples. One illustrates the independence of the optimal bidding strategy from the parameter α when, after the fact, losers know how much the winner paid. In the second example, non-winners do not know what the winner

paid, and then including regret in the utility function hurts the bid-taker's expected revenue at equilibrium. We conclude with a brief summary and overview.

The Model:

This section defines our model and establishes our main result. In particular, we consider an independent, but not necessarily privately known, values model with risk neutral bidders. Roughly speaking, in such a model, if losers learn the winners price, then bidders' optimal strategies do not depend on the relative weight given to expected profit versus expected regret.

Specifically, consider a sealed bid auction for a single object. The bid-taker has a known reservation price of r ; hereafter, we treat this reservation price as if it were simply a bid by the bid-taker. A known number, possibly random, of risk neutral bidders submit sealed bids for the object. The highest bidder wins the object, and pays an amount equal to his bid.

Bidders obtain private information--information beyond that which all bidders know--by observing the outcome of random variables with a known joint distribution. Specifically, look at the problem from the viewpoint of a particular bidder i . Let x denote the outcome of X observed by i , and let $\underline{y}$ denote the vector of outcomes of $\underline{Y}$ observed by other bidders. For later convenience, let y_j denote the j th component of $\underline{y}$, and let y^* (respectively Y^*) denote the largest component of $\underline{y}$ (respectively $\underline{Y}$). Assume that the random variable $Y^*|X = x$ has a probability density function for each possible x .

Bidder i can estimate his value from the object conditional on having observed x . However, this estimated value need not be independent of others' observations $\underline{y}$. Therefore, let $v(x, \underline{y})$ denote the expected value of the object to i conditional on x and $\underline{y}$, where the expectation is over any components of the true state of nature that would still be uncertain even if i were to have observed both x and $\underline{y}$. Assume that $v(x, \underline{y})$ is bounded.

After observing their respective component of $(x, \underline{y})$, each bidder bids. Assume that for some reason--perhaps the auction has a symmetric Nash equilibrium--each of the other bidders bids as if he were simply substituting the outcome he observed into some function $b(\cdot)$; this $b(\cdot)$ may, but need not, correspond to a Nash equilibrium. (Although we do not require the model to be symmetric, the existence of a symmetric equilibrium arises most naturally from symmetric models.) We assume that the function $b(\cdot)$ increases monotonically and is continuously differentiable; Milgrom and Weber (1982) establish appropriate conditions so that Nash equilibrium strategies satisfy our assumptions.

Eventually, at some time after he submitted his bid, i learns something additional about the auction. In particular, i learns whether or not he won the object. We presume that i learns the winner's price. In addition, i might learn something about how others bid. Since we assume that each of the other bidders follows the monotonic bidding strategy $b(\cdot)$, any information about others' bids reveals specific information about $\underline{y}$. Thus, without loss of generality, let $w(\underline{y}, x)$ denote what i learns about $\underline{y}$ after the fact given that he observed x ; $w(\cdot)$ may be vector valued. For example, if when i loses,

he still eventually learns how much the winner paid for the object, but nothing else about how others bid, then $w(\underline{y}, x)$ is informationless-- for example, $w(\underline{y}, x)$ equals some constant. Alternatively, if i always learns how its strongest competitor bid, then we might define $w(\underline{y}, x) = y^*$.

Although i does not know $\underline{y}$, or even w , at the time of bidding, he does know something about how they are related to the x that he observed. Specifically, let $F(y^*|x)$, $G(w|x, y^*)$, and $H(\underline{y}|x, y^*, w)$ denote the conditional cumulative probability distribution functions of y^* , w , and $\underline{y}$. Assume that i knows these functions, or more precisely, that he acts as if he were maximizing his expected utility with y^* , w , and $\underline{y}$ distributed according to these functions.

If i had known w before he bid, then he might have preferred to bid differently. For example, if i knew how much the highest other bidder would bid, then if i had thought of bidding more than this, he would probably want to reduce his bid to just a hair above the highest other bid; this would result in essentially zero money being left on the table. Thus, we define the regret suffered by i if he wins with a bid of β to be $\beta - b(y^*)$. Similarly, i might suffer regret when he loses if his expected value for the object exceeds his estimate of what the winner paid for it. Thus, we define the regret that i suffers when he loses to be $\int v(x, \underline{y}) dH(\underline{y}|x, y^*, w) - b(y^*)$ if this difference is positive, and zero otherwise. Then at the time of bidding, i has an expected total regret of

$$R(x, \beta) = \int_{y^* \leq b^{-1}(\beta)} (\beta - b(y^*)) dF(y^*|x) + \int_{y^* \geq b^{-1}(\beta)} \int_w^{\max\{0, \int v(x, \underline{y}) dH(\underline{y}|x, y^*, w) - b(y^*)\}} dG(\underline{w}|x, y^*) dF(y^*|x) \text{ from bidding } \beta \text{ after having observed } x. \text{ In}$$

addition, at the time of bidding, i has an expected profit of $\Pi(x, \beta) =$

$$\int_{y^* \leq b^{-1}(\beta)} v(x, \underline{y}) dH(\underline{y} | x, y^*, w) dG(w | x, y^*) - \beta dF(y^* | x).$$

Of course, by definition, i chooses to bid that value of β which maximizes his expected utility from the auction. We have already assumed that the expected utility to i would be a weighted combination of expected profit and expected regret. Thus, the expected utility to i from bidding β after observing x , knowing that others will follow the strategy $b(\cdot)$, and that he will eventually see $w(\cdot)$ evaluated at whatever $\underline{y}$ others observed in addition to learning whether or not he won the object and the winner's price may be written as follows:

$$U(x, \beta) = \alpha \Pi(x, \beta) + (1-\alpha) R(x, \beta)$$

Note that i suffers regret on losing the object only if his expected value for the object (given whatever information he now has) exceeds his estimate (also given whatever information he now has) of the price that the object went for. The case of i just barely losing the object will be of crucial importance. So, define the following condition: For any given x , α , and optimal bid β by i given x and α ,

$$\int_{\underline{y}} v(x, \underline{y}) dH(\underline{y} | x, Y^* = b^{-1}(\beta)) \geq \beta \quad \forall w: dG(w | x, Y^* = b^{-1}(\beta)) > 0 \quad (*)$$

In words, if i just barely misses winning the object, then he always expects the object to have been worth at least as much as he bid for it. This condition will hold in many, if not most real world auctions.

Theorem: If 1) at the end of the auction, a risk neutral bidder i learns both whether or not he won the single object being sold, and the

winner's price; 2) when i just barely misses winning the object, he always expects the object to have been worth at least as much as he bid for it; and 3) all other bidders bid as if they were following the same monotonically increasing, continuously differentiable bidding strategy; then i's optimal bid is independent of the relative weights given to profit and regret in his utility function.

Proof: For a bid $\beta \geq r$ to be optimal, it must satisfy the first order condition $\frac{dU(x, \beta)}{d\beta} = 0$. But, for our model, this condition becomes

$$\begin{aligned} & \alpha \left[\frac{1}{b'(b^{-1}(\beta))} \left(\int \int v(x, \underline{y}) dH(\underline{y}|x, Y^*=b^{-1}(\beta), w) dG(w|x, Y^*=b^{-1}(\beta)) - \beta \right) dF(b^{-1}(\beta)|x) - F(b^{-1}(\beta)|x) \right] \\ & = (1-\alpha) [0 + F(b^{-1}(\beta)|x)] \\ & - \frac{1}{b'(b^{-1}(\beta))} \int \max\{0, \int v(x, y) dH(\underline{y}|x, Y^*=b^{-1}(\beta), w) - \beta\} dG(w|x, Y^*=b^{-1}(\beta)) dF(b^{-1}(\beta)|x). \end{aligned}$$

Now, using condition (*) and then combining like terms reduces the first order condition to

$$\begin{aligned} & \int \int v(x, y) dH(\underline{y}|x, Y^*=b^{-1}(\beta), w) dG(w|x, Y^*=b^{-1}(\beta)) dF(b^{-1}(\beta)|x) \\ & = \beta dF(b^{-1}(\beta)|x) + b'(b^{-1}(\beta)) F(b^{-1}(\beta)|x) \end{aligned}$$

In addition, bidding $\beta < r$ results in zero profit and positive expected regret of i's expected value for the object exceeds the reservation price r . Thus i should bid less than r if and only if his expected value for the object conditional on winning with a bid $\beta=r$ is less than r . Thus both the condition for when i should bid at least r , and the condition for what i should bid if he bids at least r are independent of the relative weight given profit versus regret.

Thus, roughly speaking, if losers learn what price the winner paid, then the optimal bidding strategies are independent of α , and any theory for expected profit maximizing bidders applies just as well if bidders actually maximize some weighted average of expected profit and expected regret as we have defined it.

Examples:

This section provides two examples illustrating the above theorem. In the first example, the stated conditions for the theorem hold, and the optimal bid for i to make having seen x is independent of α . However, in the second example, if i loses, he does not learn the winner's price; this violates the conditions assumed by the theorem. In this example, i 's optimal bid does depend on α . In fact, giving equal weighting to profit and regret results in a strictly lower expected equilibrium price than if bidders simply maximized expected profit without regard to regret.

We consider two examples with independent private values. Each of the n risk neutral bidders knows his own value for the object. The values are independent samples from the uniform distribution on the unit interval. In both examples, the reservation price equals zero, and the statistic $w(\underline{Y})$ is informationless.

In the first example, each of the other bidders bids $(n-1)/n$ of his actual value. At the end of the auction, i knows the winner's price, and therefore can infer the winner's value for the object. We show that i can do no better than to always bid $(n-1)/n$ of his own value for the object. (Thus, we have found a Nash equilibrium. The fact that this Nash equilibrium is independent of α illustrates the theorem.)

To determine i 's optimal bid, start with three observations. First, the other bidders will never bid greater than $(n-1)/n$. If i bids at least $(n-1)/n$ then he always wins. Bidding greater than $(n-1)/n$ increases the price without increasing the probability of winning. Thus, i should never bid greater than $(n-1)/n$. Second, i should never bid greater than his own value x ; bidding greater than his own value may result in i winning the object at a price greater than his own value, an outcome to be avoided, and which can be avoided by bidding no more than x . (Note that therefore, condition (*) will hold in any private values example.) Third, a negative bid has essentially the same effect as a bid of zero. Thus we need only consider nonnegative bids $\beta \leq \min\{x, (n-1)/n\}$.

Now, calculate the expected utility $U(x, \beta)$ to i from bidding β when his value is x . For $0 \leq \beta \leq \min\{x, (n-1)/n\}$, this expected utility equals

$$\alpha \int_{y=0}^{\frac{n\beta}{n-1}} (x-\beta)(n-1)y^{n-2} dy - (1-\alpha) \left[\int_{y=0}^{\frac{n\beta}{n-1}} \left(\beta - \frac{ny}{n-1}\right)(n-1)y^{n-2} dy + \int_{y=\frac{n\beta}{n-1}}^{\frac{nx}{n-1}} \left(x - \frac{ny}{n-1}\right)(n-1)y^{n-2} dy \right]$$

where y denotes the largest of the other bidders' values. Differentiating this expression with respect to β , setting the result equal to zero, and solving for β yields $\beta = (n-1)x/n$ independent of α . Thus, i 's optimal bid is independent of α . Furthermore, since i 's optimal strategy is to follow the same strategy as followed by the other bidders, we have a symmetric Nash equilibrium, and the equilibrium strategy is independent of α .

On the other hand, to illustrate the crucial role played by the assumption that i learns the winner's price, now consider a second

example. In this second example, losers only learn that they lost. If $\alpha = 1$, then we get the same Nash equilibrium as before--whether or not losers know the winner's price does not effect the Nash equilibrium (or any other optimal strategy for that matter) if i bids solely to maximize expected profit.

However, in contrast to the case $\alpha = 1$, now consider the case $\alpha = 1/2$ and $n = 2$. Then, for i to have no better alternative than to follow the strategy $b(\cdot)$ when the other bidder follows the strategy $b(\cdot)$, we need the following first order condition to hold:

$$\frac{d}{d\beta} \left[\int_{y=0}^{b^{-1}(\beta)} (x-b)dy - \int_{y=0}^{b^{-1}(\beta)} (\beta-b(y))dy \right] \Big|_{\beta=b(x)} = 0 \text{ when } \int_{y=x}^1 (x-b(y))dy < 0$$

and

$$\frac{d}{d\beta} \left[\int_{y=0}^{b^{-1}(\beta)} (x-b)dy - \left[\int_{y=0}^{b^{-1}(\beta)} (\beta-b(y))dy + \int_{y=b^{-1}(\beta)}^1 (x-b(y))dy \right] \right] \Big|_{\beta=b(x)} = 0$$

$$\text{when } \int_{y=x}^1 (x-b(y))dy \geq 0.$$

In addition, $b(x)$ must be continuous x, and, if i bids r, then his value x for the object must have been zero; specifically, $b(0) = 0$.

Consider the following candidate for the equilibrium strategy:

$$b(x) = \begin{cases} \frac{x}{3} & \text{for } x \leq x^* \\ \frac{x}{2} - \frac{x^{*2}}{6x} & \text{for } x \geq x^*, \end{cases}$$

$$\text{where } \int_{y=x^*}^1 (x^*-b(y))dy = 0$$

and so x^* equals approximately 0.3148.

To verify that this satisfies the necessary first order condition, observe that

$$\int_{y=x}^1 (x - (\frac{y}{2} - \frac{x^{*2}}{6y})) dy$$

is nonnegative for $x^* \leq x \leq 1$, and for this range of x , $b(x) = x/2 - C_1/x$ satisfies the appropriate first order condition. In addition,

$$\int_{y=x}^{x^*} (x - \frac{y}{3}) dy + \int_{y=x^*}^1 (x - (\frac{y}{2} - \frac{x^{*2}}{6y})) dy$$

is nonpositive for $0 \leq x \leq x^*$, and for this range of x , $b(x) = x/3 + C_2$ satisfies the appropriate first order condition. Finally, $b(0) = 0$ and $b(x)$ is continuous at $x = x^*$ if $C_1 = x^{*2}/6$ and $C_2 = 0$. In fact, the stated strategy is unique (symmetric) equilibrium strategy for this example.

Notice that this equilibrium differs from that when $\alpha = 1$ (where $b(x) = (n-1)x/n = x/2$ for $n = 2$ was the unique equilibrium strategy). Thus, the result of the theorem depends critically on the loser knowing the winner's price. Furthermore, at least in this example, if the loser does not learn the winner's price, then the bid-taker's expected revenue suffers; the $b(x)$ defined above is strictly less than $x/2$ for all x greater than zero, and therefore the two bidders bid less when $\alpha = 1/2$ than when $\alpha = 1$.

Summary:

We consider a model of auctions and competitive bidding for a single object in which the bidders, although risk neutral, consider regret in

addition to profit when deciding how to bid. As we defined it, regret comes from two sources--the winner paying more for the object than the second highest bidder bid, and a loser possibly having a value for the object in excess of what the winner paid for it. Then, if losers learn what the winner's price, a bidder's expected utility maximizing bidding strategy is independent of the relative weight given to regret versus profit; the first example illustrates this theorem.

A second example illustrates the importance of losers learning the winner's price. In the second example, losers do not learn the winner's price and the Nash equilibrium bidding strategy depends on the relative weight given regrets versus profits. In fact, the bid-taker's expected equilibrium drops as the relative weight on regret increases.

Throughout the paper, we gave the regret suffered by a loser equal weight to that suffered by a winner. In practice, however, the regret suffered by a winner--the money left on the table--may have a larger weight than the opportunity losses of a loser. But, if we start with equal weightings on both types of regret, and then move toward more weight on winner's regret, a bidder's optimal bid decreases--he becomes relatively more concerned about overbidding than before and therefore should decrease the expected amount of overbid somewhat at the expense of somewhat increased, but less heavily weighted, regret from underbidding and losing when his value exceeds the winner's price. Of course, the amount by which the bidder's optimal bid decreases depends not only on the relative weight on the two types of regret, but also on the relative weight on regret versus profit. In short, the heavier the weight on regret from money left on the table compared to loser's regret or

compared to expected profit when winner's regret has a higher weight than losers' regret, the lower the bidder's optimal bid.

This effect on a bidder's expected utility maximizing bid of the relative weights on the two types of regret provides the intuition behind the second example. In particular, by giving losers less information about the winner's price, losers must calculate their regret based on the expected winner's price. However, even if the expected winner's price exceeds a loser's value, the winner's price might have some positive probability of exceeding the loser's value. Thus, without specific information about the winner's price, the loser may have no regret even though in the same situation, if he were told the winner's price, the loser would have positive regret with positive probability--therefore a positive expected regret. As a result, the less information a loser has about the winner's price, the less regret he suffers on average when losing. So, in effect, not knowing the winner's price reduces the relative weight on loser's regret from the weight given loser's regret when losers know the winner's price. Thus, losers not knowing the winner's price decreases the optimum bid, as illustrated in the second example.

In summary, for sealed-bid, first-price auctions of a single object to risk neutral bidders, a bidder's expected utility maximizing bidding strategy does not depend on the relative weight given to regret versus profit so long as losers learn the winner's price and both types of regret have equal weight in the bidder's utility function. Making the winner's regret over money left on the table more important than loser's regret over lost opportunities decreases the optimal bids. Not telling


losers the winner's price in effect puts less weight on loser's regrets relative to winner's regret and therefore also decreases the optimal bid. In either case, the amount of decrease in the optimal bid depends on the relative weight given regret versus profit.

Thus, this establishes the importance of losers knowing the winner's price and of both types of regret having equal importance if the theory derived from models with expected profit maximizing bidders is to be applied to cases where regret enters into bidders' utility functions in the way modelled in this paper. Roughly speaking, the theory for expected profit maximizing bidders applies just as well to bidders concerned about regret if and only if losers learn what price the winner paid and if bidders place equal weight on both types of regret. Therefore, the previous models' exclusion of regret from the utility functions affects the applicability of the theory for risk neutral bidders in first-price sealed-bid auctions only in certain, now clearly delineated, cases.

References:

- Bell, D. E., "Regret in Decision Making Under Uncertainty," Operations Research, Vol. 30, No. 5 (1982), pp. 961-981.
- Engelbrecht-Wiggans, R., "Auctions and Bidding Models: A Survey," Management Science, Vol. 26, No. 2 (1980), pp. 119-142.
- Engelbrecht-Wiggans, R., "Auctions with Bidders Subjected to Budget Constraints," to appear in International Journal of Game Theory (1987).
- Milgrom, P. R. and R. J. Weber, "A Theory of Auctions and Competitive Bidding," Econometrica, Vol. 50 (1983), pp. 1089-1122.
- Myerson, R. B., "Optimal Auctions," Mathematics of Operations Research, Vol. 6, No. 1 (1981), pp. 58-73.
- Vickrey, W., "Counterspeculation, Auctions, and Competitive Bidding," Journal of Finance, Vol. 16, No. 1 (1961), pp. 8-37.

UNIVERSITY OF ILLINOIS-URBANA



3 0112 060296040