INFORMATION AND THE DIVERGENCE BETWEEN WILLINGNESS-TO-ACCEPT AND WILLINGNESS-TO-PAY

Charles D. Kolstad
and
Rolando M. Guzman*

Revised: December 1998

*Department of Economics and Bren School of Environmental Science and Management, University of California, Santa Barbara, CA 93106-9210 and National Planning Office, Dominican Republic. Research supported in part by NSF Grant SBR-9496303. Valuable comments from Robert Deacon, Charles Plott, Jason Shogren and two anonymous referees are much appreciated.

ABSTRACT

There is considerable empirical and experimental evidence that there is a divergence between willingness-to-accept compensation to give up a good and willingness-to-pay to obtain a good. This divergence persists even when the good in question in small relative to income, a result in apparent conflict with standard economic theory. This paper develops a theoretical bidding model with costly information acquisition to explain this divergence. The model generates a gap between offers to sell and bids to buy consistent with the experimental results. We argue that the model does a better job of explaining empirical and experimental data than either of the two commonly invoked theoretical explanations: the endowment effect and the substitution effect.

Information and the Divergence Between Willingness-to-Accept and Willingness-to-Pay

I. INTRODUCTION

There has been a significant debate over the past two decades about an apparent empirical anomaly regarding willingness-to-accept compensation (WTA) and willingness-to-pay (WTP) for a good. The received wisdom (e.g., [21]) is that if a good has value that is small relative to income, your WTP to obtain that good should be very close to your WTA to give up that good. Yet numerous empirical studies, including many laboratory experiments, do not support this equality.[1] These studies generally show *announced* WTA significantly in excess of *announced* WTP. The issue is of very real empirical significance since such measures of value are often used as the basis for major public policy decisions.

The literature has been divided in explaining this anomaly. One strand of the literature involves explaining why a person's true WTA might differ from her true WTP. Another strand of the literature starts with the assumption that there is no difference between an individual's true WTP and WTA, only a difference between the individual's announced WTP and WTA.

There have been two primary theoretical arguments for why an individual's true WTP differs from her true WTA. One involves a modification of standard utility theory, positing that the value of a loss is different from the value of a gain, even if the loss or gain is small relative to income. Put another way, there is an aversion to loss in general. This theory has been put forward by Kahneman and Tversky [14] in the context of loss aversion and by Thaler [20] who coins the term "endowment effect." The meaning is that when a good becomes part of your endowment, the value you place on it increases. Thus WTA will be larger than WTP. The endowment effect has been explored in the context of experimental markets by Kahneman et al [13]. Their basic evidence is an experiment involving the random division of a group of individuals in half with half receiving some good that has value; they then to allow trading to take place. Equivalence of WTA and WTP should imply that approximately half of the goods were improperly allocated and should be traded. In fact, far less than half are traded, suggesting WTA > WTP.

The other theoretical explanation for WTA > WTP comes from Hanemann [10], although his argument is confined to relatively unique non-market goods, for which substitutes are not readily available. He gives the example of Yosemite National Park: your willingness to pay to move Yosemite into your consumption bundle may be far less than your willingness to accept compensation for removing it. He demonstrates that if there is no market substitute for the unique good,[2] then it is not possible to compensate for its removal; thus WTA is infinite whereas WTP is finite, since WTP is bounded by income. Alternatively, if there is a low-cost market good that is a perfect substitute, WTA = WTP. Another way of interpreting this is that if a non-market good lacks substitutes and is thus highly valued, the implicit income of the consumer is much higher with the commodity than without and thus WTA will not equal WTP due to an income effect. Shogren et al [19] present experimental results that they argue support the substitution effect while refuting the endowment effect. We discuss these arguments in more detail later in this paper.

Several other explanations for the WTP-WTA divergence have been articulated, though more loosely. Most of these arguments focus on the difference between *announced* WTP and WTA rather than differences in *actual* WTP and WTA. One quite straightforward argument pertains to repeated trials or auctions of the same item. In a repeated common value auction with uncertainty on the common value, and in which each agent buys or sells only one item once, each repetition provides information on the common value. Thus there is informational value in observing the outcomes of future repetitions. However, for an individual agent, this information value drops to zero once a transaction is consummated. This provides an incentive to understate WTP or overstate WTA. In both cases, the value of waiting is the option value and is the amount that WTA or WTP are biased.

A related model is presented by Crocker and Shogren [4]. Their basic assumption is that the bidder is uncertain about her valuation but learns about that value through consumption. Thus in a dynamic setting, winning a good in an auction provides utility plus information on value.[3]

Another explanation has been put forward by Hoehn and Randall [12] who argue that researching one's preferences to find one's true valuation on an item is a time consuming process and that if time is limited, agents will stop before finding the true value. They focus on an individual's problem of expenditure minimization (e.g., in the context of the expenditure function). They view this process as starting with a set of bundles that yield a certain level of utility. Expenditure minimization involves evaluating the cost of each bundle, slowly approaching the bundle, which costs the least. If a new good is introduced, maximum willingness-to-pay is the difference between income and the "current" calculation of the minimum expenditure. Thus analogously to gradually computing a lower and lower expenditure necessary to achieve a given level of utility, the calculated maximum willingness-to-pay will gradually rise. Thus a premature termination of this process leads to an understatement of WTP. Their argument is not totally convincing since there are many other paths to the minimum expenditure associated with a given level of utility, including searching outside the feasible set.

This paper approaches the issue of divergence between these values from the point of view of auction theory. Specifically, we view an agent's articulation of a WTA or a WTP as analogous to the process of forming a bid in an auction. If the bidder is unsure of his or her value of the commodity then we show how WTA will be biased upwards and WTP downwards relative to the true valuation or market clearing price.[4]

We also explicitly model the process of costly information acquisition. In the limit, if the potential payoff is small from participating in an auction relative to information costs, bidding will typically occur in an environment of ignorance with implications for announced bids. On the other hand, if information costs are low relative to payoffs, bidders will become informed before bidding. The better informed the bidders, the smaller the divergence between WTA and WTP.

Acquiring better information can take many forms. If a contractor is bidding on the construction of a building, resources can be expended to reduce uncertainty and more narrowly define his or her cost of construction. However, in the cases most cited in the literature on WTA vs. WTP, uncertainty is usually of one of two forms. In one case there may be some unknown common value associated with the good being auctioned (a common value auction). An alternative case is that there is a private value associated with the utility that the good provides; thus the value will vary from participant to participant (an independent private values auction). In the first but not the second case, the value others place on the good provides information; thus in repeated auctions for the same good, information is provided by each auction's outcome. In both cases, the bidders may expend effort (at a "cost" in terms of utility) to narrow the uncertainty on their valuation or as Cummings and Harrison [5] phrase it, "research their preferences." Although such effort may not involve the actual expenditure of money to reduce uncertainty, its effects are identical: without a sufficient potential payoff, bidders will not attempt to narrow uncertainty. Plott [17] postulates a similar process, referring to it as the "discovered preference hypothesis."

Our results show that if the payoff is small relative to the cost of information, bidders will tend to understate WTP and overstate WTA; in particular, the average value (over all bidders) of WTP will be less that the average WTA. Further, as the information costs increase (and the number of informed bidders decreases), the gap between WTP and WTA for informed bidders widens. Although we are unable to reach conclusions about how information costs effect the gap between mean WTP and WTA over *all* bidders, our theoretical model is qualitatively consistent with the experimental results of Shogren et al [19] without resorting to the substitution effect.

Our results more directly explain the Kahneman et al [13] results without invoking the endowment effect. As mentioned earlier, they examine the expected number of trades when identical commodities are randomly distributed within a group (in an experimental setting). They attribute a lower than expected number of trades to the "Endowment Effect." In our model, as the information costs increase relative to potential payoffs, the number of trades decreases.

Thus we explain the experimental results on the WTP-WTA divergence without resorting to loss aversion or the endowment effect. The next section of the paper presents our simple model of bidding behavior with costly information acquisition. We then turn to an interpretation of some of the experimental results that have been published elsewhere, arguing that our model provides a better explanation for the results than either the substitution effect or the endowment effect. We close with conclusions.

II A SIMPLE MODEL OF BIDDING BEHAVIOR

The basic problem we consider is an auction where bidders have private values on the items being auctioned but the bidders know those values imperfectly. Furthermore, bidders may narrow the uncertainty on those values at some cost. Ideally, bidders could choose any amount of information, from none, to a little to a lot with larger amounts of information costing more. Unfortunately, as Milgrom [16] points out, auction theory can become excessively complex very quickly; a complex process of information acquisition is not tractable.

Thus we confine ourselves to the simplest structure which still captures the essence of the behavior of interest. Bidders are either completely informed or uninformed. Rather than integrate suppliers and consumers into the same market, we separately consider the bid formation process on the part of buyers and on the part of sellers. On the buying side, we consider a one-shot sealed-bid, first-price auction of a single item. Bidders have private values drawn from a common distribution. Bidders know the common distribution but do not know their own private value (nor anyone else's). They may learn their true value by paying a price, known to them, and drawn from a common distribution.

Analogously, we consider the closely related case where all bidders own one unit of the item which they are offering for sale to a single buyer, who will accept the lowest bid, purchasing one unit. The bids in the first case correspond to announced WTP and in the second case to announced WTA.

Most empirical papers that find a gap between WTP and WTA fit this mold of separately considering sellers and buyers in a first-price auction (e.g., [2]). In contrast, experimental markets often involve a double auction with buyers and sellers interacting and sometimes a second-price auction (e.g., [19]). These two situations are obviously related, though not identical. Our results apply to the first case of two distinct single sided first-price auctions. We will argue that our qualitative results are relevant to the double-sided auction. The results do not apply to a second-price auction.

A. A Private Value, Costly Information Auction

Consider first the case of N risk neutral bidders participating in a first-price sealed-bid auction, bidding on a single object.[5] Assume that the object has some value, v_i, for each bidder, i, but the value v_i is unknown, even to the agent i herself. In fact, the valuations of the object are random variables with identical continuous density functions, $f(\cdot)$, with support $[V_l,V_u]$ and with expected value $\bar{v}$. This distribution is common knowledge. Further, we assume that each agent can buy perfect knowledge of her own valuation at a cost c_i. Each agent knows her own

c_i but not that of other agents. All c_i are drawn from the same distribution $g(\cdot)$, known to all. This is our private value, costly information auction.

We will assume that an agent can follow one of two alternative strategies. First, she might purchase perfect information at price c_i in order to discover the true value that the object represents to herself. Then, based on her findings, she would form an optimal bid so that her expected profit is maximized conditioned on the new information she has received. On the other hand, this agent might decide to save c_i and just use the expectation of v_i as an estimate of her (unknown) true valuation. In this case, she would form an optimal bid such that her expected profit is maximized conditioned on her assumptions and (trivial) information.

In these circumstances, the agent should calculate the expected gain that information might provide to her, before making any decision regarding whether or not to buy information. She must evaluate all the possible outcomes she could face in case she decides to be informed. These outcomes obviously depend on the behavior of the other bidders.

We now consider that each bidder assumes that a proportion p of bidders choose not to be informed about their valuations and consequently are using as their private value, the common expectation, $\bar{v}$. Therefore, if a bidder pays c_i and then discovers that her valuation is v, her optimal bidding strategy would depend only on v and the assumed p.[6] Thus the bid function would have the form $b_i = B(p,v)$, and her expected gross profit, given that particular valuation, would be p(p,v). Since *ex ante,* v is a random variable, p(p,v) is also random, and the value of the information is given by its expected value. Define expected profit as ? (p) = E(p(p,v)), where the expectation is over the v and is based on the density function $f(\cdot)$. This is the expected profit <u>before</u> information is acquired.

On the other hand, if the agent refuses to buy information, she will find it optimal to bid just $\bar{v}$, given her conjecture that $pN-1>0$ other bidders are behaving in an identical way. The argument is similar to that for marginal cost pricing in a Bertrand equilibrium.[7]

Obviously, the decision of whether or not to acquire information depends simply on the relation between the value of the information--i.e., the positive amount of expected gross profit that information provides--and the cost of information acquisition. For a particular p, the relation ? (p) > c_i only depends on c_i. If c_i is too high, agent i with refuse to be informed. Therefore a rational expectations equilibrium is given by a value p* such that the actual proportion of bidders for which ? (p*) < c_i is equal to p* itself.[8]

<u>B. Buyers</u>

Given the assumptions above, it is easy to see that for any buyer who bids b while her valuation is v = $\bar{v}$, profits are given by

$$p(p,v) = (v - b)\, F\,[\,B^{-1}(b)\,]^{[(1-p)N-1]} \qquad (1)$$

where p is the proportion of uninformed bidders. Eqn. (1) is straightforward to interpret. On the right hand side of eqn. (1), the first term is the surplus while the second term is the probability of winning. The term in braces is the probability that another informed bidder's value is less than v (ie, $B^{-1}(b)$). The probability that all the other (1-p)N-1 informed bidders have values less than v involves raising the term in brackets to the (1-p)N-1 power. Guzman and Kolstad [8] show that this yields a bid function (for informed bidders) of the form

$$B(p,v) = v - \frac{\int_{\underline{v}}^{v} F(\Theta)^{[(1-p)N-1]}\, d\Theta}{F(v)^{[(1-p)N-1]}} \qquad (2)$$

Clearly there is an incentive to shade the bid downwards from the true valuation. If the bidder were to bid her true valuation, there would be no surplus and thus nothing to be gained by winning. By lowering the bid, the surplus increases but of course the probability of winning decreases. There is some happy medium between increasing surplus and decreasing the probability of winning, resulting in the bid function given in eqn. 2. This is a standard result in first price sealed bid auctions: there is an incentive to understate value [15].

The bids given by eqn (2) are for the case where v = $\bar{v}$. In the cases where an informed bidder's actual value is equal to or less than the expected value, there is some ambiguity on the actual bid. Making a bid less than the expected value has a zero probability of winning; making a bid greater than one's valuation yields negative surplus. Thus any bid between v and $\underline{v}$ will yield exactly the same expected surplus: zero. In essence, there is a multiplicity of equilibria for valuations less than $\bar{v}$. If they were to bid $\bar{v}$ - e, e > 0, that would also be a Nash equilibrium. We assume that informed bidders bid their actual valuation when there is a zero probability of winning.

To decide whether or not to acquire information, an agent must compute her expected profit from becoming informed. To do this, she substitutes eqn. (2) for b_i in eqn. (1). The expected profit before information is acquired is the expected value of eqn (1), taken over the distribution of possible values for v:

$$\Pi(p) = \int_{\bar{v}}^{V_u} \left[\int_{\bar{v}}^{\Phi} (F(\Theta))^{[(1-p)N-1]} \, d\Theta \right] f(\Phi) \, d\Phi \qquad (3)$$

She only buys information if ? (p) > c_i. Thus the proportion of informed bidders is the same as the proportion of bidders with information costs less than ? (p). It follows that rationality of expectations requires p* to satisfy

$$1 - G(? (p^*)) = p^* \quad (4)$$

This completes our description of our *private value, costly information rational expectations buying* equilibrium: bids satisfy eqn (2) where the proportion of uninformed bidders is given by a simultaneous solution of eqns. (3) and (4). Note that agents whose true value exceeds v, will shade their bids downward from actual valuations.

This can be seen in Figure 1 where we show a possible distribution function[9] for valuations (line A). Also shown on the graph is a similar distribution function for the (1-p) informed bidders (line B) and the resulting bid function (solid line C). Note the shading of the bids.

The mean bid can easily be computed from eqn (2), though eliminating p is more difficult:

$$\overline{WTP} = \bar{v} - (1-p) \int_{\underline{v}}^{V_u} \int_{\underline{v}}^{j} \left[\frac{F(\Theta)}{F(j)} \right]^{[(1-p)N-1]} d\Theta \; f(j)\, dj \qquad (5)$$

We see right away that the mean bid (WTP) is less than the mean valuation, which yields our first result, which is a known result for first-price auctions (see [15]):

Prop. 1: In a private value, costly information rational expectations buying equilibrium, the mean bid for a single item being auctioned is less than the mean valuation (over all bidders) of that item.

It is straightforward to develop some comparative statics results. Note first of all that as the proportion of uninformed bidders increases, the shading of the informed bids increases, holding the distribution of valuations and information costs constant. This can be seen from inspection of (2), noting that F(?)/F(v) = 1.

<u>Lemma 1:</u> In a private value, costly information rational expectations buying equilibrium, an increase in the proportion of informed bidders, holding constant the distribution of values, results in an increase in the bid for any particular valuation in excess of the mean valuation.

Since in equilibrium, the proportion of uninformed bidders is endogenous, the only things that could change the proportion are changes in the distributions of values or costs. Consider first a reduction in information costs.

<u>Prop. 2:</u> In a private value, costly information rational expectations buying equilibrium, a mean-reducing shift in the distribution of information costs results in an increase in the proportion of informed bidders and thus an increase in the mean informed bid.

The interpretation here is that as it becomes less costly to acquire information, more bidders will become informed. When there are more informed bidders, each informed bidder will raise her bid to increase the chances of winning the auction.

A mean-preserving shrinking of the distribution of values is more difficult to evaluate; in fact the effect appears to be ambiguous. But what is clear is that as uncertainty on values disappears, so that all bidders know and share the same valuation, the bids all converge on that valuation.

We should also point out that Lemma 1 and Proposition 2 indicate what happens to the bids of informed bidders. To determine the effect of changes in p on mean WTP, eqn. (5) must be differentiated with respect to p and the sign of ? WTP/? p evaluated. The sign of this derivative is ambiguous.[10]

C. Sellers

We now consider the very similar situation of the N bidders selling the single item being auctioned; ie, each bidder possesses the item but only one will be purchased, at the lowest bid.

Rather than repeat all of the development of the previous section, we will highlight the changes appropriate for sellers. It is easy to see that for any seller who bids b while his

$$\Pi(p,v) = (b - v)\,[\,1 - F(\,B^{-1}(\,b\,))^{[(1-p)N-1]}\,] \qquad (6)$$

valuation is v = $\bar{v}$, profits are given by

where p is the proportion of uninformed bidders. This yields an offer function of the form

$$O(p,v) = v + \frac{\int_{v}^{\bar{v}} [1 - F(\Theta)^{[(1-p)N-1]}]\, d\Theta}{[1 - F(v)^{[(1-p)N-1]}]} \qquad (7)$$

Clearly there is some shading of the offer upwards from the true valuation. If the bidder were to offer his true valuation, there would be no surplus and thus nothing to be gained by winning. By raising the bid, the surplus increases but of course the probability of winning decreases. There is some happy medium between increasing surplus and decreasing the probability of winning, resulting in the offer function given in eqn. 7.

The offers given in eqn (7) are for the case where v < $\bar{v}$. In the cases where a bidder's actual value is equal to or greater than the expected value, there is some ambiguity on the actual offer. Making an offer greater than the expected value has a zero probability of winning; thus all bids greater than this yield zero surplus. In essence, there is a multiplicity of equilibria for valuations greater than $\bar{v}$. As in the case of bidding to buy an item, we assume that sellers offer their actual valuation (or $\bar{v}$ if they are uninformed) when there is a zero probability of winning.

We can thus consider the expected profit before information is acquired as the expected

$$\Pi(p) = \int_{V_l}^{\bar{v}} \left[\int_{f}^{\bar{v}} [1 - F(\Theta)^{[(1-p)N-1]}] \, d\Theta \right] f(f) \, d f \qquad (8)$$

value of eqn (6), taken over the distribution of possible values for V:

We can redefine an auction equilibrium, using similar language as before. A *private value, costly information rational expectations selling* equilibrium is a set of offers satisfying eqn (7) where the proportion of uninformed bidders is given by a simultaneous solution of eqns. (8) and (5). In this case, we obtain a shading of offers upward from actual valuations for those informed bidders whose actual valuations are less than the mean valuation.

This can also be seen in Figure 2 (analogous to Fig. 1) where we show a possible distribution function for valuations (line A). Also shown on the graph is the valuation of the (1-p) informed bidders (line B). Only informed bidders whose valuation is less than $\bar{v}$ have a positive probability of winning. Their offers are shown by the heavy line C in the Figure. Note the upward shading of the offers. They are re-ordered from lowest to highest and plotted with the offers of the others in line D, to form an aggregate WTA function.

The mean offer can easily be computed from eqn (7):

$$\overline{WTA} = \bar{v} + (1-p)\int_{V_l}^{\bar{v}} \left[\int_{f}^{\bar{v}} \left\{ \frac{1 - F(\Theta)}{1 - F(f)} \right\}^{[(1-p)N-1]} d\Theta \right] f(f)\, d f \qquad (9)$$

We see right away that the mean offer (WTA) is greater than the mean valuation:

Prop. 3: In a (selling) private value, costly information rational expectations bidding equilibrium, the mean offer for the item being sold through auction is greater than the mean valuation of that item.

Several comparative statics results can be developed, analogously to the case of the buyers' equilibrium.

Lemma 2: In a (selling) private value, costly information rational expectations bidding equilibrium, an increase in the proportion of informed bidders, holding constant the distribution of values, results in a decrease in the offer for any particular valuation which is less than the mean valuation. This results in a decrease in the expected informed offer, where the expectation is taken over all informed bidders.

Since in equilibrium, the proportion of uninformed bidders is endogenous, the only things that could change the proportion are changes in the distributions of values or costs. Consider first a reduction in information costs.

Prop. 4: In a (selling) private value, costly information rational expectations bidding equilibrium, a mean-reducing shift in the distribution of information costs results in an increase in the proportion of informed bidders and thus a decrease in the mean informed offer.

The interpretation here is that as it becomes less costly to acquire information, more bidders will become informed. When there are more informed bidders, each bidder will lower his offer slightly to increase the chances of winning the auction.

A mean-preserving shrinking of the distribution of values is more difficult to evaluate; in fact the effect appears to be ambiguous. But what is clear is that as uncertainty on values disappears, so that all bidders know and share the same valuation, clearly the bids all converge on that valuation.

D. WTA vs. WTP

We have separately considered the strategies of buyers and sellers. We will now compare WTA and WTP, though not by integrating buyers and sellers into a double-sided auction. Instead we consider two separate identical sets of agents. One set is a set of buyers; the other is a set of sellers. Each set is the same size with identical valuations and information costs. Buyers and sellers form bids as discussed above.

Although short of a true double sided auction, this framework can yield useful insights on WTA vs. WTP issues. The two results we present here follow from the lemmas and propositions above. One (which directly follows from the existing literature) is that for any individual, announced WTP = WTA. Consequently, the mean {WTP} = mean {WTA} over the population of bidders. Except in cases where all bidders are uninformed, the inequality is strict. The second result is that as information costs decrease, the gap between *informed* WTP and WTA shrinks.

The first result follows directly from Prop. 1 and 2:

Prop. 5. In a private value, costly information rational expectations bidding equilibrium involving buyers and sellers, for any individual bidder, the announced WTP is less than or equal to the announced WTA. For informed bidders, the inequality is strict. Consequently the mean WTP over the population of bidders is less than the mean WTA.

Thus we have the basic result that theoretically, one expects announced willingness-to-pay to be less than announced willingness-to-accept without resorting to a modification of utility theory to include loss aversion or an endowment effect and without invoking a lack of close substitutes to yield the difference. This result follows directly from the theory of Nash bidding in a first-price auction.

Information costs play a crucial role in determining the gap between WTP and WTA:

Prop. 6. In a private value, costly information rational expectations bidding equilibrium involving buyers and sellers, a mean-reducing shift in the distribution of information costs will increase the number of informed bidders and as a consequence decrease the gap between mean informed WTA and WTP.

Proposition 6 indicates that the more costly information is, the larger the expected gap between informed WTA and WTP. Thus for goods whose value is difficult to assess, one would expect a large gap. This is in fact what often occurs for goods that are unfamiliar, such as environmental goods. This is a new result that hinges on the change in the number of informed bidders as information costs change.

E. Potential Trades

Part of the motivation for this paper is the experimental result that in double-sided experimental auctions, WTA appears to exceed WTP. One particular manifestation of this is the experiment of Kahneman et al [13]: when half the bidders are randomly endowed with a good, the fraction of the bidding population who would voluntarily trade based on announced

WTP and WTA is generally less than 50%.

Although we are unable to examine a true double-sided auction within our theoretical framework, we can develop a related concept within our two separate single-sided markets. For our two identical groups of buyers and sellers, define the number of *potential trades* as the maximum number of pairs of bidders (one buyer, one seller) where the buyer's bid exceeds the seller's offer. Because we are not explicitly considering a double-sided auction, this definition of potential trades may not correspond to the number of trades in an actual double-sided auction.

Prop. 7. In a private value, costly information rational expectations equilibrium, with information costs such that there are informed and uninformed bidders, a mean-increasing shift in the distribution of information costs results in a decrease in the number of potential trades.

The proof of this proposition is largely an application of propositions 1 and 2. As long as the proportion of bidders is an interior point (ie, neither 0 nor 1), then an increase in information costs results in more bidders being uninformed. Uninformed bidders bid the expected value of the item and will not enter into trade. Only buyers who bid more than $\bar{v}$ will potentially buy and only sellers who bid less than $\bar{v}$ potentially sell. Thus the lower the proportion of informed bidders, the lower the number of potential trades. Only with all bidders informed will the number of potential trades approach N/4.

III A REEXAMINATION OF EXPERIMENTAL RESULTS ON WTP VS. WTA

We now return to the discussion earlier in the paper on the evidence in the literature on the divergence between WTP and WTA.

Kahneman et al (1990) argue for the endowment effect and as evidence cite an experiment involving a group of 44 individuals, where 22 were randomly given Cornell

University coffee mugs. They conduct sequential double-sided auctions. Logically, approximately half of the people who received the mugs value them at less than the average in the group; similarly, half the people who did not receive them should value them at more than the average. Thus they argue that 11 trades would be expected but they only observe an average of 3 trades over 4 repetitions of the experiment. They conclude that WTA must be significantly higher than WTP and attribute the difference to the endowment effect.

Shogren et al [19] refute the Kahneman et al [13] results by repeating the experiment with Iowa State University coffee mugs but running ten repetitions of the market. They show an initial wide divergence between WTA and WTP but as the auction is repeated, the gap shrinks, largely disappearing after 4 repetitions. While these results are somewhat at odds with the results of Kahneman et al [13], they do tend to refute the existence of an endowment effect.

How well does our information theoretic model explain these results?[11] Because we are unable to model a complete double-sided auction, we use the concept of "potential trades" as a proxy for what might happen in a double-sided auction. Certainly the small number of initial trades found by Kahneman et al [13] is consistent with our results. The item in question is a market good where the main uncertainty regards its market price. Even if a bidder is in love with a mug and willing to pay $100, he or she knows it can be easily purchased at the market price and thus should not fetch a price in excess of that market price. With repeated trials the bidders are obtaining information on what the group views the market price as being. And after repeated trials, the uncertainty is largely eliminated and WTP=WTA. Thus the Shogren et al [19] results are qualitatively consistent with our model.

A second set of evidence is provided by Shogren et al [19]. They set up an experiment whereby subjects are given a choice between fast food with a normal risk of food poisoning and fast food that has received additional screening for pathogens. The authors show a substantial divergence between WTA and WTP, with the gap persisting over as many as 20 repetitions. They suggest that this is evidence of Hanemann's substitution effect driving a wedge between WTP and WTA. While one's life (or even good health) clearly has no substitutes, it is not clear that a very low probability of sickness has no close substitute. People often substitute money, and not very much of it, for small changes in low probability risks to life and health.

How well does our model explain these theoretical results? In this case, the value of the good is private to a bidder and only based on the bidder's value on health (roughly). Thus knowing how your neighbor values her health provides no information to you on how you value your health. Your own valuation can only be discerned by "researching your preferences." Thus repeated trials would not provide any information and thus we would not expect to see any change from trial to trial. Furthermore, the good in question (health) is difficult to value (high information costs), and the payoff from better information is potentially low, since we are quite accustomed to food with a "normal" risk. Consequently, one would expect a fairly substantial divergence between mean informed WTA and WTP. Although we are unable to reach conclusions about the size of the gap between mean WTA and WTP (over all the bidders), the fact that the experimental results of Shogren et al [19] do show a significant gap which remains unchanged with repetitions of the auction, is consistent with an information-theoretic explanation for the WTP-WTA gap.

IV CONCLUSIONS

In this paper we have examined the long-standing debate of why there is a difference between willingness-to-accept compensation to give up a good and willingness-to-pay to obtain a good when that good appears to be of small value relative to the buyer's or seller's income. This apparent disparity has arisen time and again in both experimental settings and actual markets.

Two theories have been put forward by others to explain the disparity. Kahneman et al

[13] modify utility theory to include an endowment effect whereby mere ownership of a good increases its value. Hanemann [10] suggests that much of the evidence for divergence involves goods which have no close market substitutes and thus conventional utility theory can be expected to generate a divergence (a substitution effect). Shogren et al [19] present experimental results testing these two theories and claim to refute the endowment effect while supporting the Hanemann substitution effect.

This paper takes a different approach. We develop a model of auction equilibrium where bidders do not know their true valuation but can expend effort or money to learn that valuation. We show in this context that a divergence between willingness-to-accept and willingness-to-pay naturally emerges and further, that the divergence increases when the cost of information acquisition is greater.

Using this model, we reinterpret the experimental results of Kahneman et al [13] and Shogren et al [19] and argue that our model does a better job of explaining the results than either the endowment effect or the substitution effect.

Two significant and one more modest caveats are appropriate here (and suggest further research). First of all, we are examining a first-price auction. We also examine an approximation of a double-sided auction through our use of the concept of "potential trades." Much of the *experimental* evidence on the disparity between WTA and WTP is based on double-sided auctions, and in some cases, second-price auctions. Empirical evidence is from a context closer to a first-price auction. Thus our model only partially addresses the empirical/experimental anomaly.

A second significant point is that learning and incomplete information are represented in the simplest possible fashion. It would be more realistic and interesting to allow players to acquire different levels of information (at cost) and then engage in bidding

V REFERENCES

1. G. Bhattacharya, The Importance of Learning in Planning with Exhaustible Resources, *in* "Essays in the Economics of Renewable Resources" (L. Mirman and D. Spulber, Eds.), Elsevier, Amsterdam (1982).

2. R.C. Bishop, T.A. Heberlein and M.J. Kealy, Contingent Valuation of Environmental Assets: Comparisons with a Simulated Market, *Nat. Res. J.*, 23, 619-33 (1983).

3. D.L. Coursey J.L. Hovis and W.D. Schulze, The Disparity Between Willingness to Accept and Willingness to Pay Measures of Value, *Quart. J. Econ.*, 102, 679-90 (1987).

4. T.D. Crocker and J.F. Shogren, Preference Learning and Contingent Valuation Methods, *in* "Environmental Policy and the Economy" (F. Dietz, F. van der Ploeg and J. van der Straaten, Eds.), Elsevier, Amsterdam (1991).

5. R.G. Cummings and G.W. Harrison, Was the Ohio Court Well Informed in its Assessment of the Accuracy of the Contingent Valuation Method?, *Natural Resources Journal* (forthcoming).

6. M. Demers, Investment Under Uncertainty, Irreversibility and the Arrival of Information over Time, *Review of Economic Studies*, 58, 333-50 (1991).

7. S.J. Grossman and J.E. Stiglitz, The Impossibility of Informationally Efficient Markets, *American Economic Review*, 70, 393-408 (1980).

8. R.M. Guzman and C.D. Kolstad, Bidding Behavior under Costly Information Acquistion: An Experimental Study, Working Paper in Economics #17-95, Department of Economics, University of California, Santa Barbara (1995).

9. R.M. Guzman and C.D. Kolstad, Auction Equilibrium with Costly Information Acquisition, Working Paper in Economics #10-95R, Department of Economics, University of California, Santa Barbara (1997).

10. W.M. Hanemann, Willingness to Pay and Willingness to Accept: How Much Can They Differ?, *Amer. Econ. Rev.*, 81, 635-47 (1991).

11. G.W. Harrison, Theory and Misbehavior of First-Price Auctions: Reply, *Amer Econ. Rev.*, 82, 1426-43 (1992).

12. J.P. Hoehn and A. Randall, A Satisfactory Benefit Cost Indicator from Contingent Valuation, *J. Env. Econ. and Mgmt.*, 14, 226-47 (1987).

13. D. Kahneman, J.L. Knetsch and R.H. Thaler, Experimental Tests of the Endowment Effect and the Coase Theorem, *J. Pol. Econ.*, 98, 1325-48 (1990).

14. D. Kahneman and A.Tversky, Prospect Theory: An Analysis of Decision under Risk, *Econometrica*, 47, 263-91 (1979).

15. R.P. McAfee and J. McMillan, Auctions and Bidding, *J. Econ. Lit.*, 25, 699-738 (1987).

16. P.R. Milgrom, "The Structure of Information in Competitive Bidding" Garland Publishing, New York, (1979).

17. C. Plott, Rational Individual Behavior in Markets and Social Choice Processes, Caltech Social Science Working Paper 862 (Dec. 1993).

18. A. Randall and J.R. Stoll, Consumer's Surplus in Commodity Space, *Amer. Econ. Rev.*, 71, 449-57 (1980).

19. J.F. Shogren, Y.S. Seung, J.H. Dermot and J.B. Kliebenstein, Resolving Differences in Willingness to Pay and Willingness to Accept, *Amer. Econ. Rev.*, 84, 255-70 (1994).

20. R. Thaler, Toward a Positive Theory of Consumer Choice, *J. Econ. Behav. and Organization*, 1, 39-60 (1980).

21. R. Willig, Consumer's Surplus Without Apology, *Amer. Econ. Rev.*, 66, 589-97 (1976).

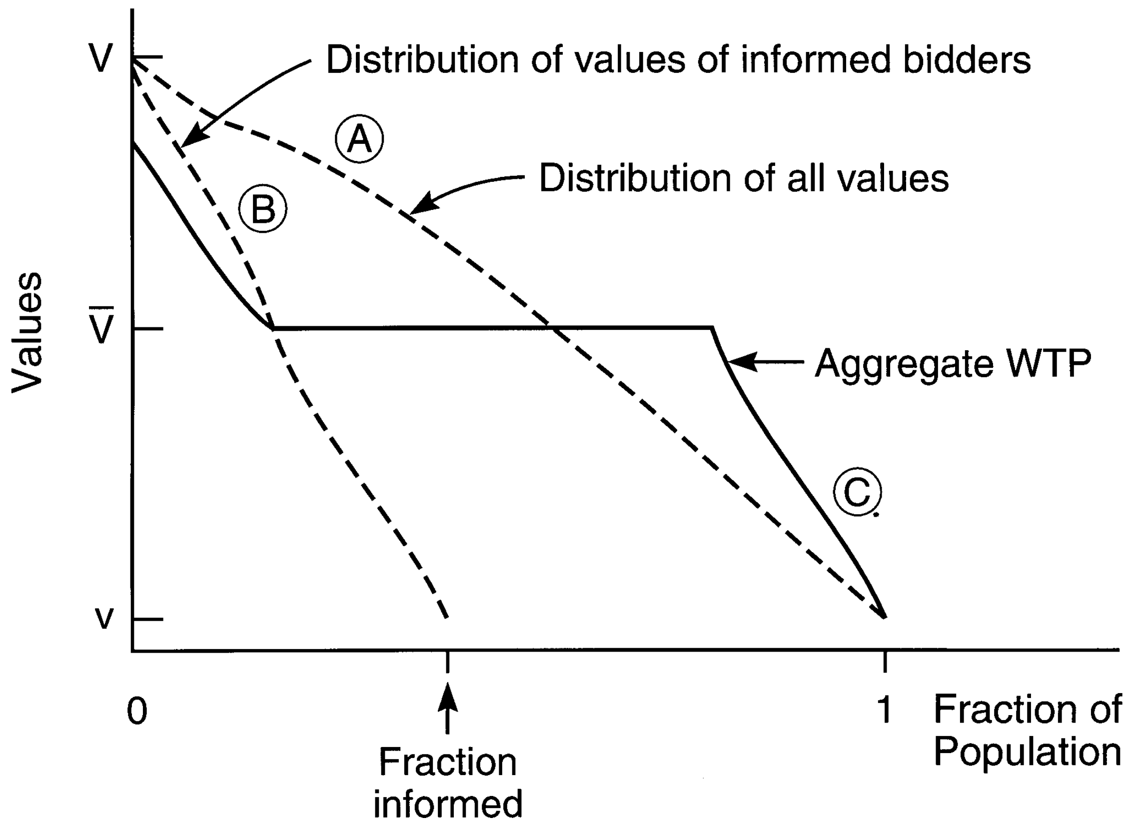

Figure 1: Hypothetical distribution of valuations and resulting aggregate willingness-to-pay function.

Note: For a given value, V_i, the distribution of values gives the fraction with valuations greater than V_i.

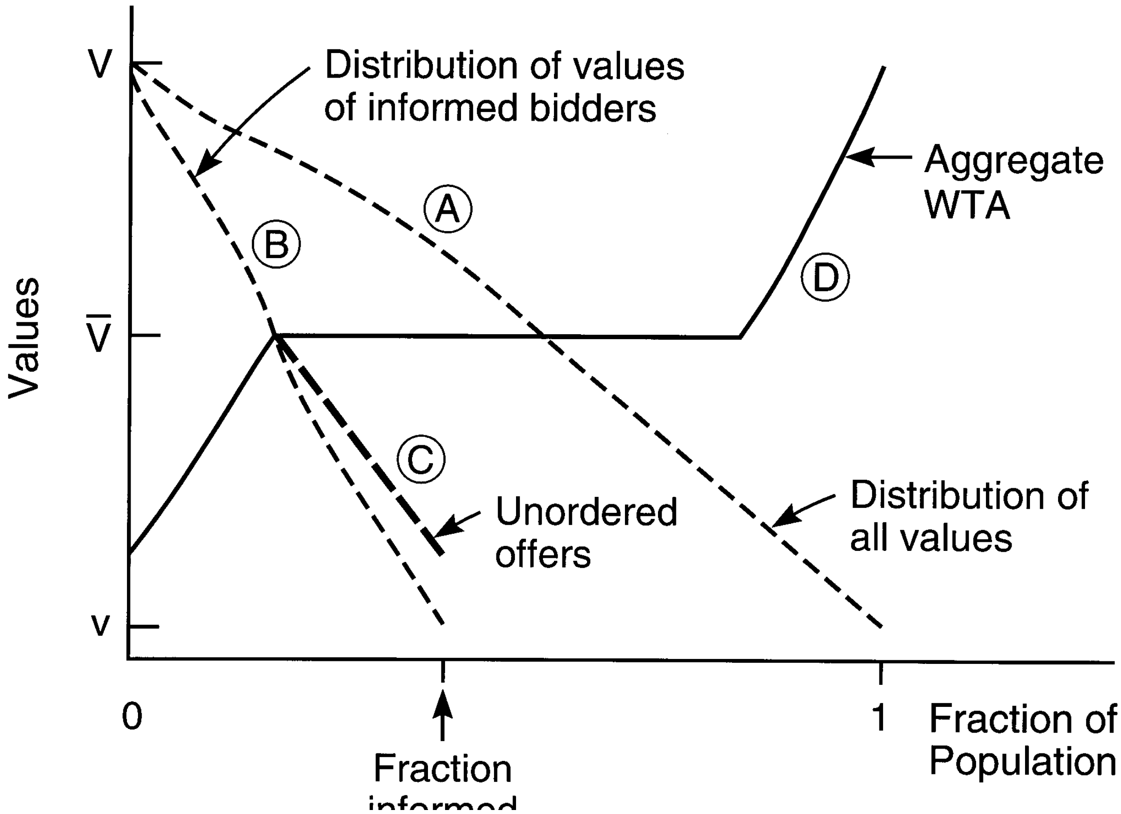

Figure 2: Hypothetical distribution of valuations and resulting aggregate willingness-to-accept compensation function.

See note on Figure 1.

ENDNOTES

[1]Shogren et al (1994) provide the most recent experimental results on this issue. Kahneman et al (1990) provide their own experimental results and review a number of other studies that have shown that WTA exceeds WTP. Both of these papers are experimental. Coursey et al (1987) use a second-price auction for a bad (sucrose octa-acetate) and argue that WTP and WTA converge in repeated trials. Harrison (1992) summarizes several of these results, pointing up flaws in the experimental design and specifically disputing the results of Coursey et al (1987). One of the most interesting non-experimental studies is by Bishop et al (1983) who compare a contingent valuation measurement (based on self-reporting of values) of WTA and WTP to a market clearing price in a constructed market for goose hunting permits.

[2]By no market substitute, Hanemann (1991) means no amount of private goods can compensate for taking away the non-market good.

[3]This result on optimal decisions with learning by experimentation has appeared in a number of other contexts. For example, Bhattacharya (1982) considers exploration for resources as providing new resources as well as information about the unknown stock. Demers (1991) considers the case of competitive firm which is uncertain about demand and which is seeking to find the optimal capital stock.

[4] This is a direct consequence of the known result that in a first price sealed bid auction, bidders will shave their private value when bidding.

[5]In a first-price sealed-bid auction, each bidder submits his bid in a private way, without observing the bids offered by his competitors. The auction is won by the bidder who

presents the highest bid (in the case of acquiring the good) or the lowest bid (in the case of selling the good).

[6]The equilibrium concept used here is not the Bayesian Nash equilibrium, where each agent uses the joint distribution on costs and values (types) to compute the optimal bid. Such an equilibrium concept is intractable (and not even appropriate) for the problem at hand. Consequently, we have adopted a simplified equilibrium concept, similar to that used by Grossman and Stiglitz (1980).

[7]Assume there are at least two uninformed bidders and applied to symmetric equilibria where all uninformed bidders act the same. Guzman and Kolstad (1995) consider other cases. Suppose there is a symmetric Nash equilibrium with all uninformed bidders bidding $ß_j < \bar{v}$. Any bidder bidding less that $\bar{b} = \max_j ß_j$ can increase profits by increasing her bid to $\bar{b} + e$ where $e > 0$ and $\bar{b} + e < \bar{v}$. This process continues until all are bidding $\bar{v}$.

[8]Here we ignore the discrete nature of p. See Guzman and Kolstad (1995) for more detail on this.

[9]The distribution function is more precisely defined as: for a given value V_i, the distribution function gives the fraction with valuations greater than V_i.

[10]Consider what happens as p decreases. More bidders become informed, increasing the mean informed bid, which is less than $\bar{v}$. This tends to increase the overall mean bid. But the number of uninformed bidders (bidding $\bar{v}$) drops. This tends to reduce the overall mean bid. The net result is ambiguous without additional structure on the problem.

[11]The information theoretic model presented here has also been tested in an

experimental market (see Guzman and Kolstad, 1995).