An Experimental Investigation of the Seller Incentives in the EPA's Emission Trading Auction

By TIMOTHY N. CASON*

The Clean Air Act requires the EPA to conduct annual auctions of emission allowances. Under the discriminative auction rules, sellers with the lowest asking prices receive the highest bids. This paper studies an inverted version of this auction in which buyers face the same incentives as sellers in the EPA auction. Consistent with theoretical predictions, buyers bid above their valuation, auction outcomes are inefficient, and increasing the number of buyers increases bids. Buyers facing human opponents compete more aggressively than the risk-neutral prediction, but bids do not differ systematically from this prediction when buyers face computerized Nash "robots." (*JEL* D44, L51, Q25)

The Clean Air Act Amendments of 1990 ("the Act") created tradable emission allowances to control nationwide sulfur dioxide pollution. In theory (W. David Montgomery, 1972), allowance trading provides flexibility in achieving air quality objectives, thereby lowering the total cost of emission reductions. To ensure the availability of the emission allowances and to provide clear price signals to the evolving allowance market, Congress instructed the U.S. Environmental Protection Agency (EPA) to conduct annual sealed-bid/sealed-offer auctions. Unfortunately, the Act states that "allowances shall be *sold on the basis of bid price*, starting with the highest-priced bid and continuing until all allowances for sale at such auction have been allocated" (*Clean Air Act Amendments of 1990* [Public Law 101-549], Sec. 416(d)(2); emphasis added). The EPA interprets this language of the Act to require a discriminative auction with a unique feature: all sellers in this market—including private firms—receive the bid price of a specific buyer. Sellers with the lowest asking prices receive the highest bids. This paper reports ten experimental sessions that demonstrate the poor performance properties of this new trading institution.

An earlier paper (Cason, 1993) models this institution as a Bayesian-Nash game of incomplete information to study its incentive properties. Although the EPA "does not believe that a systematic reduction in minimum [asking] prices will occur" (*Federal Register*, 17 December 1991, p. 65596) as a result of the low-offer-receives-high-bid rules, the model demonstrates that these rules can create a strong incentive for sellers to misrepresent and underreveal their true costs of emission control. The misrepresentation incentive arises because lower asking prices increase the probability that a seller trades with high-bidding buyers; therefore, lower asking prices only indirectly reduce the expected price received. Asking prices that are below the true cost of emission control, however, are risky because they may cause sellers to trade at a loss. This risk limits the amount of profitable misrepresentation. In the symmetric equilibrium with differentiable offer strategies, the benefits outweigh the risks of misrepresentation so that on balance risk-neutral

*Department of Economics, University of Southern California, Los Angeles, CA 90089-0253. Financial support was provided by the Randolph Haynes and Dora Haynes Foundation. James Cox, Dan Levin, participants at the Fall 1992 Economic Science Association and Winter 1993 American Economic Association meetings, and especially an anonymous referee provided helpful comments on earlier drafts. I retain the responsibility for any remaining errors. Zagros Madjd-Sadjadi provided very capable research assistance.

sellers have the incentive to offer units at prices substantially below their cost of emission control. Increasing the number of sellers does not eliminate this misrepresentation, and misrepresentation can intensify with increased seller competition if the number of buyers remains fixed. Therefore, the EPA's rules have the potential to generate significantly biased price signals and reduce the efficiency of the sulfur dioxide allowance auction.

The experiment reported here supports the main comparative-static implications of the Nash model but rejects the precise risk-neutral equilibrium. The experiment implements a one-sided version of the auction to focus on the seller incentives and to control for the strategic behavior of buyers.[1] One unit is auctioned each period. Furthermore, the experiment inverts the institution so that a set of buyers faces incentives that are strategically equivalent to those facing sellers in the EPA auction.[2] In this simplified environment, the high-bidding buyer wins the auction and pays an asking price (*not* her bid price) randomly drawn by nature if her bid price exceeds the asking price. The results are consistent with the main theoretical predictions: Buyers bid above their true valuation for the object, auction outcomes are inefficient, and an increase in the number of buyers increases bids. However, buyers facing other human opponents compete *more* aggressively than predicted by the risk-neutral Bayesian-Nash model. Risk aversion does not explain this result. Bids do not differ systematically from the risk-neutral Nash equilibrium in a treatment in which buyers face computerized Nash "robots," however, so the overbidding against human opponents may be due to strategic responses to a subset of aggressive, high-bidding subjects.

[1]James C. Cox et al. (1984) explore the theoretical and behavioral properties of the multiple-unit discriminative auction.

[2]Selling units to buyer subjects seems more natural to explain in instructions and allows the results to be directly compared to other experimental auction results with different pricing rules. It also simplifies the explanation of the (Nash) robot strategies used in the "robot-opponent" treatment.

The remainder of the paper is organized as follows. Section I briefly summarizes the rules of the EPA auction institution. Section II presents a simple one-sided version of the Bayesian-Nash model of this auction in which buyers' bids determine their probability of winning and the range of possible bid prices they pay. The end of Section II contains the experimental hypotheses. Section III summarizes the experimental design, and Section IV contains the results. Conclusions are presented in Section V. Experiment instructions are available from the author upon request.

I. A Brief Summary of the EPA Auction

The centerpiece of the acid-rain control program in the Clean Air Act Amendments of 1990 is a system of tradable emission allowances. An allowance authorizes the emission of up to one ton of sulfur dioxide, and the total annual emissions of each "affected" unit (most units in operation prior to passage of the Act) must be less than or equal to the number of allowances held for that unit. Allowances may be transferred to and from any affected unit or person. Consequently, each unit may meet its sulfur dioxide limitation in the most efficient means possible, either by choosing the most cost-effective emission control technology or buying emission allowances from units with lower emission control costs.

To maintain a cap on total emissions, new units (most units beginning operation after passage of the Act) must obtain allowances from existing allowance-holders or through the EPA auctions and sales programs. EPA calls once-a-year auctions no later than March 31 of each year. EPA holds two separate auctions at this annual call: a *spot* auction for allowances allocated for the current year and an *advance* auction for allowances effective in seven years. The Act allows any person holding allowances to sell them in the auctions held by the EPA but requires that the (reserve) allowances in EPA's auction subaccount be sold before any private offerings. The number of EPA subaccount allowances in each of the two annual auctions varies from 150,000 to 250,000 between 1993 and 1999, but is set at

200,000 after 1999 (i.e., in "Phase II"). This amounts to 2.24 percent of the total allowances allocated each year.

The EPA auction is a new variant of the "call-market" institution used on many organized exchanges. In the call-market institution, potential buyers submit (sealed) bids specifying their maximum willingness to pay for the commodity, and potential sellers submit (sealed) offers (or "asks") indicating the minimum they are willing to accept in exchange for the commodity. (One can think of these bids and asks as "limit orders," which may specify different price limits for different units. The EPA auction subaccount allowances are sold with a minimum asking price of zero.) The market is "called," and trades are executed at a known, preannounced time using a previously specified rule for determining the price of each trade. Although these auctions are two-sided—with both bids to buy and offers to sell—the EPA interprets the language of the Act "to require that allowances be sold to successful bidders at the price of their respective bids (also referred to as a discriminative approach)" (*Federal Register*, 23 May 1991, p. 23746). The EPA determines the prices as follows (*Federal Register*, 23 May 1991, p. 23746):[3]

> All bids to the auctions will be ranked from highest to lowest on the basis of bid price. EPA will allocate and sell all the allowances in the auction subaccount on the basis of this ranking: when all such allowances are sold, EPA will match contributed allowances offered for sale with any remaining bids. Specifically, EPA proposes to match the offer to sell that stipulates the lowest minimum price with the highest remaining bid. This matching process will continue in ascending order of specified minimum price until all bids are awarded or allowances are consumed, or until EPA can no longer match bids with allowances because sellers have set their minimum price higher than any remaining bids.

Allowances sold in this matching scheme receive the bid price of a specific buyer, and *lower* asking prices increase trading priority and can thus lead to *higher* received prices. The EPA is aware that "such a system may provide incentives for holders of allowances to specify lower minimum prices for allowances than they would be willing to accept in order to be matched to higher bids" (*Federal Register*, 23 May 1991, p. 23746). It is worth noting that more standard uniform-price call auctions (such as those studied by Mark A. Satterthwaite and Steven R. Williams [1989] and Aldo Rustichini et al. [1994]) have much better theoretical incentive properties because only the marginal traders affect price.

II. A Theoretical Bayesian-Nash Model

This section presents a simplified Bayesian-Nash model of the auction that captures the salient strategic incentives induced by the EPA's low-offer-receives-high-bid rules. Buyers compete to buy a single unit, and the single nonstrategic seller (the experimenter) submits an offer price from a known distribution. The highest bidder pays the offer price, *not* her bid. Thus, this model and the experiment examine the (strategically equivalent) inverse of the problem facing sellers in the EPA auction. For a more general multiple-unit characterization expressed in seller terminology, see Cason (1993). The framework and assumptions employed here are similar to the existing auction literature based on the seminal insight of William Vickrey (1961) (see e.g., Milton Harris and Artur Raviv, 1981; John G. Riley and William F. Samuelson, 1981; Cox et al., 1982a, b, 1984, 1988).

One unit is offered at an asking price of c, where c is drawn independently (each period) from a known probability distribution function $\Phi(\cdot)$ with density $\phi(\cdot)$ whose

[3] The final rules published in the *Federal Register* on 17 December 1991 do not modify this proposal. Proceeds of sales from the auction subaccount are distributed to units that had allowances withheld from their initial allocation, so this is a revenue-neutral auction.

support is the interval $[0, \bar{c}]$. Each of $N > 1$ buyers submits a bid to buy this single unit, and the buyer with the highest bid may win the unit and pay the asking price c.[4] The random asking price is unknown when the buyers enter their bids.[5] Rank the N bids in decreasing order, $b_1 \geq b_2 \geq \ldots \geq b_N$. The purpose of the bids is only to rank the buyers. The high-bidding buyer wins the unit and pays the asking price c only if her bid $b_1 \geq c$. Note that the winning buyer acquires a lottery: should she win, she will pay an asking price that is a random variable distributed on the interval $[0, b_1]$. Let v_i denote the valuation of the unit for buyer i (this allowance valuation corresponds to the marginal emission-compliance cost of buyer i), where $i = 1, \ldots, N$. Assume that each v_i is drawn independently from a known probability distribution function $H(\cdot)$ with density $h(\cdot)$ whose support is the interval $[0, \bar{v}]$. If buyer i buys the unit at the price c, she gains the monetary surplus $\{v_i - c\}$. If buyer i does not buy a unit, which can happen if either (a) any other buyer $j \neq i$ bids $b_j > b_i$ or (b) the asking price c is greater than b_i, then her monetary gain is zero.

Assume that buyer i expects each of her rivals to bid according to the bid function

$$(1) \qquad b_k = b(v_k) \quad \forall k \neq i$$

which is nondecreasing with well-defined first and second derivatives on $[0, \bar{v}]$. Differentiability requires that $\bar{b} = b(\bar{v}) \leq \bar{c}$, where $\bar{b}$ is the maximum bid that is associated with the highest value draw.[6] Let $v(b_k)$ denote the inverse of the bid function (1):

$$(2) \qquad v_k = v(b_k).$$

The probability that bid b_i is highest is the probability that $N-1$ rivals of buyer i bid less than b_i. [This is simply the $(N-1)$th order statistic from a sample of size N. In the symmetric equilibrium, all $N-1$ rivals bid less than b_i if all $N-1$ drawn values are less than v_i.] Denote this probability $G[v(b)]$, which is given by

$$(3) \qquad G[v(b)] = [H(v(b))]^{N-1}.$$

Notice that the probability that the bid b_i is greater than the cost draw c is $\Phi(b_i)$. Finally, denote the expected utility of the bid b_i conditional upon its being the highest and greater than c as $U[b_i | b_i = \max(b_1, b_2, \ldots, b_N), c \leq b_i, v_i]$, or simply $U[b_i|\text{win}]$. Normalize the utility function so that the utility from losing the auction is zero. With this notation, the overall expected utility of a bid b_i is

$$(4) \qquad EU(b_i) = G[v(b_i)]\Phi(b_i)U[b_i|\text{win}].$$

The bid strategy is a function that maps buyer i's private value $v_i \in [0, \bar{v}]$ into the bid she submits $b_i \in \mathbb{R}_+$ given her beliefs described above. I solve for the symmetric Bayesian-Nash equilibrium bid function $b_i = b(v_i)$ by maximizing $EU(b_i)$ with respect to the bid b_i, taking the behavior of the other $N-1$ buyers as given by $b_k = b(v_k)$, $k \neq i$. The solution is a symmetric Nash equilibrium only if the optimum $b_i^* = b(v_i)$; that is, each buyer i's strategy $b_i = b(v_i)$ is the best response strategy to $N-1$ buyers playing $b_k = b(v_k)$. The first-order condition for a maximum $EU'(b_i) = 0$ implies the following nonlinear first-order differential equation for the inverse bid function:

$$(5) \qquad v'(b_i) = -\frac{G(v(b_i))[U(b_i|\text{win})\phi(b_i) + \Phi(b_i)U'(b_i|\text{win})]}{G'(v(b_i))\Phi(b_i)U(b_i|\text{win})}.$$

[4] The theoretical derivation ignores the possibility of tied bids since they are zero-probability events when drawn from continuous probability density functions. In the experiment, ties were resolved by randomization.

[5] The assumption of a random asking price is appropriate in this context but is a deviation from the standard auction literature. If buyers knew the asking price (and they do not in the EPA auction), the trivial optimal bid strategy would be to bid $+\infty$ if the asking price is below emission-compliance cost, and to bid $-\infty$ otherwise.

[6] If $\bar{b} > \bar{c}$, the derivation of the optimal bid function would be complicated by the possibility that a *range* of high value draws would map into the same optimal bid.

This differential equation and the initial condition $v(0)=0$ characterize the symmetric Nash-equilibrium bid function.[7]

Buyer values in the experiment are drawn from the uniform density $h(v)=1/\bar{v}$, and the asking price c is drawn from the uniform density $\phi(c)=1/\bar{c}$. The former condition simplifies the order statistic $G[v(b)]$ since $H(v)=v/\bar{v}$, and the latter condition simplifies $U[b_i|\text{win}]$. In particular, if all rivals bid according to the inverse bid function $v(b)$, then the probability that the bid b_i is the highest bid is

$$G[v(b_i)] = \left[\frac{v(b_i)}{\bar{v}}\right]^{N-1} \tag{6}$$

and the probability that the asking price draw c is less than b_i is

$$\Phi(b_i) = b_i/\bar{c}. \tag{7}$$

Now suppose that buyers are risk-neutral, so their objective is to maximize expected payoff. In this case simply write the utility of income $u(y)$ as y. This implies that the expected utility conditional upon winning the auction is

$$EU(b_i|\text{win}) = v_i - \frac{b_i}{2} \tag{8}$$

because the conditional expected value of the asking price is $b_i/2$. (In other words, the expected price paid conditional on winning is one-half the bid.) Substituting these expressions into equation (5) and simplifying results in

$$v'(b_i) = \frac{2v(b_i)[b_i - v(b_i)]}{(N-1)b_i[2v(b_i)-b_i]}. \tag{9}$$

Using the initial condition $v(0)=0$, it is straightforward to verify that this differential equation has the following linear solution:

$$v(b_i) = \frac{N+1}{2N} b_i \tag{10}$$

so the symmetric, risk-neutral Nash-equilibrium bid function is

$$b(v_i) = \frac{2N}{N+1} v_i. \tag{11}$$

Since $2N \geq N+1$, the optimal bid is *above* the buyer's value v_i. For the $N=3$ and $N=6$ cases of the experiment, the optimal bid is 50 percent and 71 percent higher than the value, respectively. In contrast, the equilibrium bid function in a standard first-price auction (where winning buyers pay their bids) implies that buyers bid *below* their values (Vickrey, 1961). (For the distribution functions considered here, the first-price risk-neutral bid function is $b_f(v_i) = [(N-1)/N]v_i$.) Moreover, when the number of offered units is fixed at one, increasing N under these rules *increases* the amount of misrepresentation: As $N \to \infty, b(v_i) \to 2v_i$. The bids converge to double the valuations because the expected price paid is always one-half the bid [see equation (8) above], and competition from increasing N forces expected profit to zero.[8] The data analyses below use this "perfect-competition" zero-profit bid function as a nonstrategic alternative model.

In the symmetric Nash equilibrium, buyers bid above their valuations, so the winning buyer may regret winning *ex post* (i.e., suffer losses) if the asking-price draw happens to be high. (The EPA requires all bids and asks to be binding; a certified check for the total bid price must accompany all bids.) *Ex post* winner regret in this context arises from different strategic considerations than the well-known "winner's curse" in common-value auctions (see e.g., John Kagel and Dan Levin, 1986). In the common-value

[7] It is straightforward to verify the second-order conditions in the uniform-distribution case considered in what follows.

[8] With the first-price bid function $b_f(v_i)$ above, it is easy to verify that increasing N leads to truthful value revelation; that is, $b_f(v_i) \to v_i$ as $N \to \infty$.

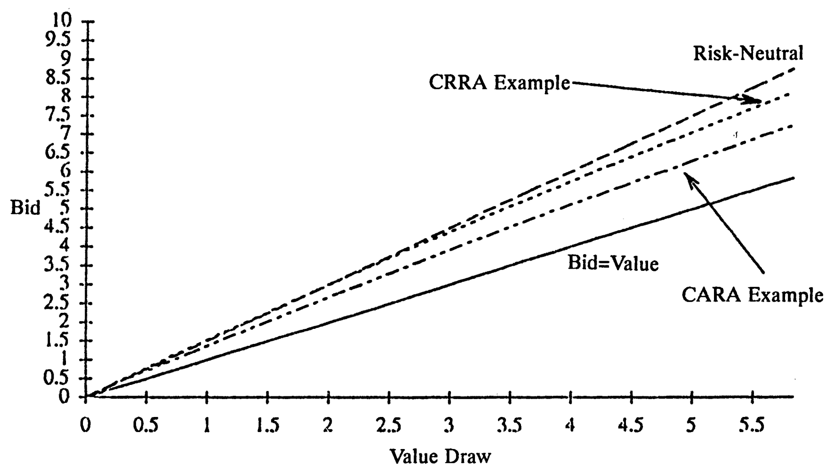

FIGURE 1. SYMMETRIC NASH EQUILIBRIUM BID FUNCTIONS FOR $N = 3$ AND EXPERIMENT PARAMETERS

auction, buyers must account for an adverse selection problem in estimating the true value of the item for sale. Buyers must discount their bids because, should they win, their signal is among the most optimistic and is therefore biased upward. The winner's curse therefore comes from judgmental failures when updating posterior expected values, and if present may cause winning bidders systematically to lose money. In contrast, the *ex post* regret suffered at times in the current model is consistent with fully rational profit maximization. The bid maximizes expected profit based on the distribution of possible cost and value draws. Although this results in losses from some cost draws, buyers earn positive profits on average.

Risk aversion explains many of the deviations from the risk-neutral Nash model in previous laboratory auction research. For example, in the first-price auction, risk aversion increases bids, lowers expected profit, and reduces value misrepresentation relative to the risk-neutral equilibrium (see e.g., Harris and Raviv, 1981; Riley and Samuelson, 1981). In the EPA auction, I have not characterized the implications of risk aversion for general risk preferences. Nevertheless, all of the utility-function parameterizations I have examined—including the commonly used constant absolute risk aversion (CARA) $[u(y) = 1 - e^{-ry}]$ and constant relative risk aversion (CRRA) $[u(y) = y^{r^*}]$ utility models—indicate the same thing: risk aversion reduces bids, increases expected profit, and reduces value misrepresentation relative to the risk-neutral equilibrium. In both the EPA auction and the first-price auction, buyers risk not winning the auction in conditions they would find profitable, which introduces upward pressure on bids relative to the risk-neutral case. However, the EPA auction has an additional downside risk: higher bids increase the probability and expected value of losses from paying more than the unit's value. This leads to lower bids than the risk-neutral prediction.[9]

[9]The third-price auction studied in Kagel and Levin (1993) has similar strategic properties. With N buyers, uniform distributions, and risk-neutrality, this institution's Nash-equilibrium bid function is

$$b_3(v_i) = [(N-1)/(N-2)]v_i.$$

As in the EPA auction, bids exceed values, and risk aversion reduces bids. Even with risk aversion, bids increase with increased N in the EPA auction. Unlike the EPA auction, however, increasing N in the third-price auction reduces bids and reduces value misrepresentation. Therefore, competition increases efficiency in the third-price auction but reduces efficiency in the EPA auction.

Figure 1 displays example bid functions for parameters of the experiments: values were drawn from [\$0.00, \$5.83] and asking prices were drawn from [\$0.00, \$9.99]. The figure presents the $N = 3$ symmetric equilibrium bid functions for the risk-neutral case ($b_i = 1.5v_i$) as well as bid functions for symmetric CRRA utility and symmetric CARA utility. (For the CRRA utility function the Arrow-Pratt constant-relative-risk-aversion parameter is $1 - r^* = 0.5$, and for the CARA utility function the Arrow-Pratt constant-absolute-risk-aversion parameter is $r = 0.5$.)[10] Notice that the risk-averse bid functions are closer to the truthful-revelation bid = value line than the risk-neutral cases. The conclusion holds for all risk-aversion parameters, not just 0.5. Thus for CRRA or CARA utility, risk aversion does not explain bids *above* the risk-neutral bid function. This conclusion is likely to hold more generally, and CARA is probably a sufficient approximation given the small money stakes in the experiment.

Finally, consider the expected profit implications of nonoptimal bidding. Define the *expected loss function* $L(b_i, b_i^*)$ as the expected loss from bidding b_i when the optimal bid is b_i^* (Daniel Friedman, 1992). If this loss function were symmetric, the payoff consequences of suboptimal behavior would suggest that errors above and below the optimal bid b_i^* are equally likely. It is straightforward to show that the loss function in this environment is approximately symmetric for small deviations (i.e., $b_i \pm$ \$0.50), but becomes asymmetric for larger deviations. For example, consider the risk-neutral equilibrium, a typical winning draw of \$5.00 and bid deviations of $\pm$\$1.00. For the $N = 3$ treatment, the expected loss is more than 40-percent greater for $b_i^* + 1$ ($b_i = \$8.50$) than for $b_i^* - 1$ ($b_i = \$6.50$). For the $N = 6$ treatment, the expected loss is more than 140-percent greater for $b_i^* + 1$ ($b_i = \$9.57$) than $b_i^* - 1$ ($b_i = \$7.57$). Because of this loss-function asymmetry, payoff-metric considerations (Glenn W. Harrison, 1989) indicate that bid "errors" are more likely to be below the optimum bid b_i^* because downward errors are less costly.

The experiment evaluates three testable hypotheses provided by the main comparative statics of the Nash equilibrium model.

Hypothesis 1: Winning bids are significantly higher for the experimental treatments with a greater number of bidders.

Hypothesis 2: All submitted bids are significantly higher for the experimental treatments with a greater number of bidders.

Hypothesis 3: Realized market trading efficiency is significantly lower for the experimental treatments with a greater number of bidders.

The Nash model also provides precise, quantitative hypotheses for all bids and efficiency under the risk-neutrality assumption, but these are easily rejected so I do not define them formally here.

The data analysis below considers a specific alternative model to the symmetric Nash model. This nonstrategic alternative is a simple rule of thumb to bid two times the drawn valuation. Like the perfect-competition (bid = value) alternative to the Nash model in the first-price auction, this alternative model is an upper bound on the "competitiveness" of buyers. The Nash model allows positive expected profits, while the perfect-competition alternative requires zero expected profits. Subjects may adopt this rule of thumb if they focus on the expected value of the cost draw conditional on winning (which, recall, is one-half the bid), ignoring the strategic aspects of the auction. I refer to this aggressive competition alternative as the *double-value* model.[11] Note that the double-value model implies that the number of competing buyers will have no impact on bidding behavior and

[10] The risk-averse utility function examples are calculated numerically using the Runge-Kutta method of numerical approximation (see e.g., William H. Press et al., 1986 Ch. 15).

[11] In the experiment, the maximum possible bid $\bar{b} = 9.99 < 2v_i$ for $v_i \geq 5.00$. Therefore, for $v_i \in [5.00, 5.83]$, the double-value model prediction is $b_i = 9.99$.

TABLE 1—SUMMARY OF EXPERIMENTAL SESSIONS

Session number	Number of buyers: First 10 periods	Number of buyers: Middle 10 periods	Number of buyers: Last 10 periods	Subject expertise	Number of buyers vs. robot opponents
E3631	3	6	3	inexperienced	0
E3632	3	6	3	inexperienced	3
E6361	6	3	6	inexperienced	1
E6362	6	3	6	inexperienced	3
E3633x	3	6	3	experienced	2
E3634x	3	6	3	experienced	1
E6363x	6	3	6	experienced	1
E6364x	6	3	6	experienced	2
All robot 1	—[a]	—[a]	—[a]	inexperienced	all 7
All robot 2x	—[a]	—[a]	—[a]	experienced	all 6

Notes: All sessions involved 30 auction periods in each human opponent session and 45 auction periods in each robot opponent session; all values were drawn from the uniform distribution over [\$0.00, \$5.83], and all costs drawn from the uniform distribution over [\$0.00, \$9.99].
[a] Varied across subjects.

efficiency. Support for Hypotheses 1–3 therefore implies a rejection of the double-value model.

III. Experimental Design and Procedures

The experiment was conducted at the University of Southern California's Experimental Economics Laboratory, which contains 17 personal computers on a local-area network. Between six and nine subjects participated in each session, all recruited from undergraduate economics classes not taught by the experimenter. Subjects were randomly assigned to computers, and instructions were read aloud by the experimenter while the subjects followed along on their own copies.[12] The subjects then continued at their own pace through two computerized examples that demonstrated the mechanics of submitting bids and illustrated the auction rules. Two practice periods then followed with no monetary payoffs. For the "human-opponent" treatment, each session contained 30 periods with monetary payoffs, and this was common knowledge. Six of the subjects participated in the human-opponent treatment each session. The computers did the profit accounting for subjects, but subjects also filled out record sheets (which contained the pre-drawn values for the entire experiment) to verify the profit calculations. Sessions lasted an average of 75 minutes, and total monetary payoffs ranged from \$4.00 to \$31.50, averaging a little more than \$20 per subject.[13]

The primary treatment variable was the number of buyers N, which varied systematically between 3 and 6 within each experiment using an ABA design (10 periods each treatment). One or two practice periods without monetary payoffs separated the treatments. In the three-buyer auction periods, each subject participated in one of two simultaneous auction markets, always competing against the same two buyers. Table 1 summarizes the 10 sessions. All sessions drew seller valuations from the discrete uniform distribution over [\$0.00, \$5.83] and drew asking prices from the discrete uniform distribution over [\$0.00, \$9.99]. (These

[12] The instructions explain how the conditional distribution of the purchase price changes as the bid changes. The computer also presented this information after subjects entered (but before they confirmed) each bid.

[13] Salient rewards, defined as the total absolute change in account balance over the course of the experiment, averaged about \$18.50 per subject.

upper endpoints were chosen because the product of 12/7 and 5.83 is 9.99, which is the maximum risk neutral bid $\bar{b}$ in the symmetric Nash equilibrium when $N=6$. For reasons described in the previous section, bids are expected to be less than or equal to the risk-neutral Nash equilibrium.) Half of the sessions used inexperienced subjects, and the other half employed these subjects in experienced sessions.[14]

As indicated in Table 1, most of the sessions included extra subjects who played against "robot" opponents. The robots played the risk-neutral Nash strategy of equation (11). These subjects served two purposes. First, they were available to replace any bankrupt subject in the human-opponent sessions in order to maintain six total buyers.[15] Fortunately, only one subject went bankrupt (in the next-to-last-period), so subjects experienced with robot opponents earlier in the same session actually made only one bid (out of the 1,440 "human-opponent" bids). Second, the robot-opponent treatment provides an additional environment to study behavior and learning against stable, known, Nash strategies. Following James M. Walker et al. (1987), a special instruction supplement describes the Nash strategies employed by the robot opponents. The subjects competing against robot opponents were able to complete each trading period more quickly than the human-opponent subjects since they did not have to wait for slower human decision-makers to make their bids; consequently, each robot-opponent treatment lasted for 45 periods with monetary payoffs. In all other respects the robot-opponent treatment was exactly the same as the human-opponent treatment. The final two sessions employed all robot opponents.[16]

IV. Experimental Results

I present results in five subsections. Section IV-A presents the overall market performance, including an analysis of the winning bids, profit, and trading efficiency. Section IV-B examines all bids pooled across all buyers, and Section IV-C considers individual bidding behavior. Section IV-D presents the behavior of buyers who faced Nash robot opponents, and Section IV-E derives best responses to the observed bid deviations from the risk-neutral Nash equilibrium.

[14]All five inexperienced sessions used a single set of value and cost draws, and all five experienced sessions used another set of value and cost draws. This procedure controls value and cost variation across replications. The number of buyers in each auction varied systematically across sessions while the value and cost draws remained constant, so this procedure also allows an exact pairing of bids with the same valuation across the buyer-number treatments. This provides powerful statistical tests for Hypothesis 1.

[15]Subjects were declared bankrupt and excused from the experiment if their current account balance ever dropped below \$5.00. This positive cutoff level ensured that when subjects bid they were able to cover their maximum possible loss. Subjects began with an initial account balance of \$15.00. This relatively large starting balance was chosen because of frequent bankruptcies in the pilot experiment.

[16]The final two all-robot opponent sessions differed slightly from the others in one way. Although all robot-opponent subjects completed a fixed number of periods (45), a referee raised the concern that because these subjects played at their own pace they could be trading off accuracy for the time required to complete the session. However, in the top eight sessions listed in Table 1, the robot-opponent subjects were explicitly told they were "backup" subjects and were required to wait for the human-opponent session to finish before being paid and excused. In order to create similar incentives for the all-robot-opponent sessions listed at the bottom of Table 1, subjects were informed that the experimenter would not pay and excuse them until a specific time. The inexperienced sessions with human opponents lasted between 82 and 106 minutes (with an average of 96 minutes), so in the inexperienced session all-robot 1 we told the subjects that they would not be paid until 80 minutes had elapsed. (Rushing to completion was apparently not an incentive for most subjects, as the average completion time was 90 minutes for this group.) The experienced sessions with human opponents lasted between 47 and 57 minutes (average = 52 minutes), so in the experienced session all-robot 2x we did not pay subjects until 50 minutes had elapsed. Under these conditions the average completion time was 55 minutes.

TABLE 2—RISK-NEUTRAL NASH THEORETICAL AND OBSERVED WINNING BIDS (HUMAN OPPONENTS): MEANS AND VARIANCES

	Six-buyer treatment		Three-buyer treatment		
Periods	(i) Theoretical Nash mean (variance)	(ii) Observed mean (variance)	(iii) Theoretical Nash mean (variance)	(iv) Observed mean (variance)	(v) Observed difference: mean (6-buyer) − mean (3-buyer)
Inexperienced:					
1–10	9.13	8.28	7.38	7.11	1.17†
	(0.58)	(1.40)	(2.22)	(2.34)	
11–20	8.37	8.35	6.55	7.54	0.81†
	(3.64)	(2.30)	(3.57)	(3.53)	
21–30	8.13	8.60	6.87	8.36	0.24
	(2.75)	(2.64)	(1.70)	(2.62)	
1–10 and 21–30, pooled	8.53	8.46	7.07	7.87	0.58
	(2.03)	(2.04)	(1.89)	(2.84)	
All 1–30, pooled	8.48	8.42	6.90	7.77	0.65*
	(2.40)	(2.08)	(2.41)	(3.04)	
Experienced:					
1–10	9.26	9.04	7.32	8.57	0.47
	(0.54)	(1.14)	(2.03)	(2.81)	
11–20	9.04	9.90	6.92	8.28	1.62*
	(0.77)	(0.03)	(1.66)	(3.13)	
21–30	8.54	9.50	6.37	8.44	1.06
	(1.23)	(0.63)	(3.51)	(5.26)	
1–10 and 21–30, pooled	8.90	9.26	6.94	8.50	0.76
	(0.96)	(0.92)	(2.73)	(4.09)	
All 1–30, pooled	8.95	9.48	6.93	8.43	1.05**
	(0.87)	(0.71)	(2.29)	(3.76)	

Notes: Statistical tests reported in column (v) are based on the random-effects model described in equation (12).
†Statistically significant at the 10-percent level.
*Statistically significant at the 5-percent level.
**Statistically significant at the 1-percent level.

A. *Winning Bids, High Bids, and Market Performance*

Table 2 presents means and variances of the winning bids in each sequence of ten auctions for the human-opponent sessions, separated by experience level and buyer-number treatment. Columns (i) and (iii) contain the risk-neutral Nash-model mean winning-bid predictions (for the particular sequence of value draws), and columns (ii) and (iv) contain the observed mean winning bids. The mean winning bid is above the risk-neutral Nash prediction in nearly all three-buyer sequences (except the initial ten periods) and in about half of the six-buyer sequences. Column (v) presents the difference in the mean six-buyer and mean three-buyer bids. Hypothesis 1 implies that this difference should be greater than 0. To test this hypothesis statistically it would be inappropriate to pool all bids and treat them as independent observations, because bids submitted by the same subject are nonindependent.

I therefore test Hypothesis 1 with the following random-effects model estimated separately for inexperienced and experienced subjects:

$$b_{it} = \sum_{k=1}^{3} \alpha_{3k} D_{3k} + \sum_{k=1}^{3} \alpha_{6k} D_{6k} + u_{it} \tag{12}$$

where b_{it} is buyer i's bid in period t, $k=1$, 2, and 3 represent the first, second, and third 10-period sequences (corresponding to the rows of Table 2), and the 3 and 6 subscripts denote three-buyer and six-buyer auctions. The D_{jk} term is a dummy variable equal to 1 for a j-buyer auction in the kth 10-period auction sequence, so the ordinary least-squares estimates $\hat{\alpha}_{jk}$ of equation (12) correspond to the means reported in columns (ii) and (iv) of Table 2. However, to account for the dependent observations across subjects I estimate equation (12) using the random-effects error specification $u_{it} = \tau_i + \varepsilon_{it}$. The term τ_i reflects the random subject-specific effect. I employ two-step generalized least-squares estimation to obtain consistent and asymptotically efficient estimates of equation (12) (Cheng Hsiao, 1986 Ch. 3). For Hypothesis 1, I test the null hypothesis $\alpha_{3k} = \alpha_{6k}$ (for $k = 1$, 2, and 3) against the alternative hypothesis $\alpha_{3k} < \alpha_{6k}$. (The pooled 20- and 30-period sequences indicated in the center and at the bottom of Table 2 employ analogous tests with different dummy variables.) Column (v) indicates that the data reject this null hypothesis of no buyer-number treatment effect at relatively weak significance levels for about half of the 10-period auction sequences; however, the data strongly reject this null when pooling across the 30 auctions of a session. I therefore conclude that the data support Hypothesis 1: increasing the number of buyers increases winning bids.

Nevertheless, one can reject the precise bid predictions of the risk-neutral Nash model for the pooled data. What is surprising is that, in all cases except for the inexperienced six-buyer treatment, this model is rejected in the direction of *more aggressive* bidding; furthermore, bids appear to increase with experience.[17] However, this test also rejects the double-value model in favor of less aggressive bidding in all cases except the experienced six-buyer treatment. At the end of the experienced sessions, high bids are about a dollar or more higher than the risk-neutral Nash-model prediction.

Table 3 presents average profit and efficiency for the auction sequences. Profit figures are average total profits across sessions (in dollars) for the indicated 10-, 20-, or 30-period sequences. Columns (i) and (iv) indicate that, pooling over all 30 periods, the Nash model predicts that this institution will extract only 69–87 percent of potential gains from exchange (efficiency numbers are reported in parentheses). However, columns (iii) and (vi) document even lower observed efficiency levels. Subjects competed away much of the available exchange surplus—typically even more than predicted by the double-value model.[18] As discussed below, the main cause of these low efficiencies is that buyers who do not have the highest-value draw often win the auction. Columns (iii) and (vi) show that average observed efficiency is greater in the three-buyer treatments than in the six-buyer treatments in four out of the six 10-period sequences. This provides weak evidence consistent with Hypothesis 3.

Both the Nash and double-value models are symmetric, and thus both imply that the buyer with the highest value submits the highest bid.[19] Therefore, ignoring tied bids, inefficiency can only occur with symmetric buyers if the winning buyer pays more than her value. In the three-buyer treatment for the values drawn in the experiment, the Nash model predicts that this high bid exceeds the cost draw and wins in about 64 percent of the periods, and in about 14 percent of the periods the winning bidder pays more than her value. According to the double-value model, the high bid exceeds the cost draw in 78 percent of the periods

[17] Formally, this statistical test is carried out using a random-effects model analogous to equation (12) with the difference between the high bid and the risk-neutral Nash model prediction as the dependent variable.

[18] Although the double-value model would appear to compete away all profits, this was not possible because bids were constrained to the cost draw range [\$0.01, \$9.99]. Therefore, the double-value model implies that bids associated with all values above \$4.99 should cluster at the upper endpoint, \$9.99. The conditional expected cost draw for bids of \$9.99 is \$4.99, so these bids had positive expected profit. Because the upper endpoint of costs was $\bar{c} = \$9.99$, allowing buyers to bid up to some point $\bar{b} > \$9.99$ would simply raise the clustering point for maximum bids without affecting the conditional cost distribution. No optimal bid would lie in the interval $[\bar{c}, \bar{b}]$, because higher bids could only increase the probability of winning and would not increase the conditional distribution of costs.

[19] Two or more buyers may submit tied bids of \$9.99 at times in the double-value model. Recall that ties are broken randomly.

TABLE 3—THEORETICAL AND OBSERVED AVERAGE PROFITS AND EFFICIENCY (HUMAN OPPONENTS)

	Six-buyer treatment			Three-buyer treatment		
Periods	(i) Theoretical Nash profit (efficiency)	(ii) Double-value theoretical profit (efficiency)	(iii) Observed average profit (efficiency)	(iv) Theoretical Nash profit (efficiency)	(v) Double-value theoretical profit (efficiency)	(vi) Observed average profit (efficiency)
Inexperienced:						
1–10	14.49	5.01	2.45	28.01	12.45	23.24
	(83.9)	(29.0)	(14.2)	(91.0)	(40.4)	(75.5)
11–20	7.87	1.37	1.91	12.80	3.81	4.14
	(68.9)	(12.0)	(16.7)	(67.2)	(20.1)	(21.8)
21–30	17.30	17.17	13.96	32.31	29.84	22.40
	(86.5)	(85.9)	(69.8)	(92.3)	(85.3)	(64.0)
1–10 and 21–30, pooled	31.79	22.18	16.41	60.32	42.29	45.64
	(85.3)	(59.5)	(44.0)	(91.7)	(64.3)	(68.9)
All 1–30, pooled:	39.66	23.55	18.32	73.12	46.10	49.78
	(81.4)	(48.4)	(37.6)	(86.3)	(54.4)	(58.7)
Experienced:						
1–10	8.92	4.26	0.52	11.28	7.24	7.95
	(73.5)	(35.1)	(4.3)	(65.3)	(41.9)	(46.0)
11–20	15.01	10.88	7.16	29.59	21.53	15.96
	(74.9)	(54.3)	(47.7)	(90.7)	(60.0)	(48.9)
21–30	10.45	7.28	7.00	27.45	8.62	−9.51
	(58.5)	(40.7)	(39.1)	(97.3)	(30.5)	(−33.7)
1–10 and 21–30, pooled	19.37	11.54	7.52	38.73	15.86	−1.56
	(64.5)	(38.4)	(25.0)	(85.1)	(34.9)	(−3.4)
All 1–30, pooled:	34.38	22.42	14.68	68.32	37.39	14.40
	(68.7)	(44.8)	(29.3)	(87.4)	(47.9)	(18.4)

Note: Efficiencies (reported in parentheses) are percentages.

but loses money in 28 percent of the periods. In the inexperienced-buyer sessions the realized frequencies lie between these two model predictions: high bids exceed the cost draw in 67 percent of the periods and lose money in 21 percent of the periods. In contrast, buyers become more aggressive in the experienced sessions: high bids exceed the cost draw in 83 percent of the periods and lose money in 34 percent of the periods. Results are similar for the six-buyer treatment, with one exception: in the experienced sessions, the frequencies of accepted bids (88 percent) and accepted bids that lose money (33 percent) are similar to the double-value model predictions (90 percent and 22 percent, respectively). Therefore, the realized efficiencies in the experienced sessions shown in Table 3 fall short of the double-value model predictions, but not because buyers are winning the auction too often compared to the double-value model. Instead, efficiency suffers primarily because buyers who do not have the highest-value draw frequently win the auction. This does not seem to be due to systematic asymmetries in the buyers' bid strategies, based on the estimated error components in the random-effects models reported above that employ the error structure $u_{it} = \tau_i + \varepsilon_{it}$. The variance in the τ_i subject-specific error term is typically one-half to one-tenth as large as the variance in the unsystematic error term ε_{it}.

B. *All Bids*

Figure 2 illustrates the observed mean bids and the Nash and double-value model mean bid predictions by buyer number treatment for the two experience levels. This figure includes all bids, including the losing bids. Observed bids are clearly lower in the three-buyer treatment than the six-buyer

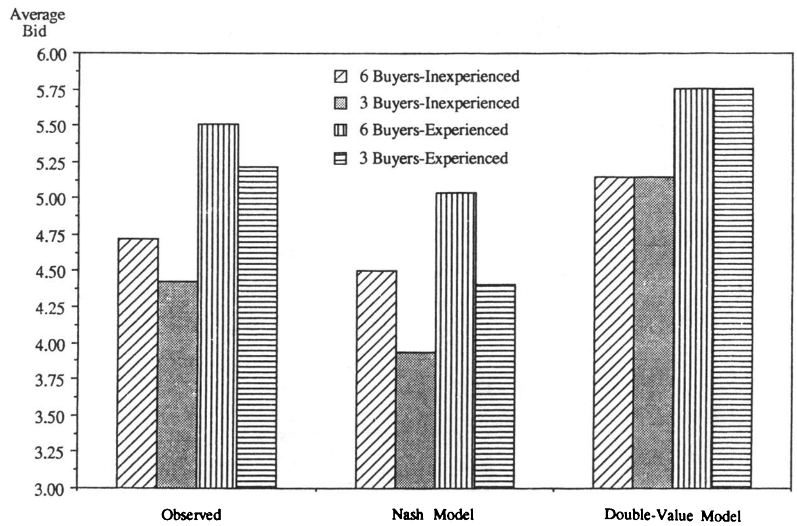

FIGURE 2. AVERAGE BIDS BY NUMBER OF BUYERS AND EXPERIENCE TREATMENTS: OBSERVED AND THEORETICAL PREDICTIONS (HUMAN OPPONENTS)

treatment, which supports Hypothesis 2. The average difference is 29 cents in the inexperienced-buyer sessions and 28 cents in the experienced-buyer sessions (paired across identical value draws), which are both significant at the 5-percent level using a random-effects model similar to equation (12). Note, however, that the mean bids for both treatments are greater than their respective risk-neutral Nash model predictions and less than the double-value model predictions. These differences are also statistically significant using the random-effects model, with one exception. In the six-buyer treatment with inexperienced subjects the test does not reject the risk-neutral Nash model.

Figure 2 indicates that average observed bids are between the two model predictions and suggests that bids become more competitive with increased experience. One can test the hypothesis that increased experience increases bids by comparing bids across experience treatments. Because values varied across sessions it is necessary to control for the drawn values. The Nash model suggests the following "value markup" equation, which can control for value variation across sessions:

$$b_{it} = \beta_1 D_1 v_{it} + \beta_2 D_2 v_{it} + u_{it} \tag{13}$$

where b_{it} is the bid of buyer i in period t, v_{it} is the value drawn by buyer i in period t, D_1 is an indicator variable equal to 1 if t is from an inexperienced session, D_2 is an indicator variable equal to 1 if t is from an experienced session, and β_1 and β_2 are "markup coefficients" to be estimated for each of the two experience levels. As in equation (12), I use a random-effects model with $u_{it} = \tau_i + \varepsilon_{it}$ and employ two-step generalized least-squares estimation. (The risk-neutral Nash model implies that $\beta_1 = \beta_2 = 1.5$ for the three-buyer treatment and that $\beta_1 = \beta_2 = 1.714$ for the six-buyer treatment.) The design suggests a logical pairing of inexperienced and experienced sessions, as each cohort of subjects that began with a "363" ("636") inexperienced session returned for an experienced 636 (363) session. This pairing of sessions for each cohort results in 4 cohorts × 2 buyer number treatments = 8 estimates of equation (13). In all 8 cases, $\hat{\beta}_2 > \hat{\beta}_1$, indicating that buyers bid more aggressively in their experienced session. The difference is significant at the 5-percent level in four of the eight cases.

C. *Individual Bidding Behavior*

Table 4 presents evidence that many individual buyers bid differently than predicted by either the risk-neutral Nash or double-value model. This table summarizes individ-

TABLE 4—NASH AND DOUBLE-VALUE HYPOTHESIS TESTS FOR INDIVIDUAL SUBJECTS FACING HUMAN OPPONENTS (PAIRED WILCOXON SIGNED RANK TESTS, 30 OBSERVATIONS PER SUBJECT)

	Nash model hypothesis		Double-value model hypothesis		
Session number and experience	Statistic too low	Statistic too high	Statistic too low	Statistic too high	Number of subjects
Inexperienced:					
E3631, E3632	4 (3)	8 (5)	9 (5)	3 (1)	12
E6361, E6362	4 (3)	8 (6)	9 (6)	3 (0)	12
All inexperienced:	8 (6)	16 (11)	18 (11)	6 (1)	24
Experienced:					
E3633x, E3634x	1 (1)	11 (7)	9 (5)	3 (2)	12
E6363x, E6364x	3 (1)	9 (5)	9 (7)	3 (3)	12
All experienced:	4 (2)	20 (12)	18 (12)	6 (5)	24
All Sessions	12 (8)	36 (23)	36 (23)	12 (6)	48

Note: The numbers in parentheses are the numbers of times the model prediction was rejected at the 5-percent level of significance (two-tailed test).

ual Wilcoxon signed rank tests for each (human-opponent) buyer that pair observed bids with the Nash and double-value model bid predictions (these tests employ 30 bid observations per subject). Overall, 36 out of 48 buyers (75 percent) bid higher than the Nash-model prediction on average, and the test rejects the risk-neutral Nash model (in either direction) at the 5-percent level for 31 subjects (65 percent). Exactly the same number of buyers (36) bid lower than the double-value model prediction on average, and the test rejects the double-value model (in either direction) at the 5-percent level for 29 subjects (60 percent). Although not shown in Table 4, results are virtually identical when considering the auctions with three buyers and the auctions with six buyers separately.

D. *Buyers Facing Nash Robot Opponents*

The symmetric Nash model correctly predicts the qualitative changes in bid behavior, but it implies lower bids than those observed in the data. Estimates of equation (13) discussed above indicate that bids tend to increase with subject experience, and subjects appear to learn to bid more aggressively as they compete against aggressive subjects. Learning is very difficult in this environment for at least three reasons. First, subjects must learn to best-respond against opponents who are also learning and revising strategies. Opponent strategies are thus often unstable. Second, subjects do not observe their opponents' strategies; instead, they only observe outcomes of their opponents' value-to-bid mapping each period. Third, opponent strategies are typically non-Nash, and the Nash strategy of equation (11) is not optimal against non-Nash opponent strategies. (Subsection E below discusses this in detail.) The robot-opponent sessions control these learning problems and test whether subjects can learn optimal (Nash) best responses to stable, known, Nash equilibrium strategies. The 26 subjects competing against Nash robots received a complete description of the robot-opponent strategy.

This treatment indicates that the risk-neutral Nash model performs better when subjects compete against computerized Nash opponents. Similar to the human-opponent bid data shown in Figure 2, bids are about 30–36 cents higher on average in the six-buyer auctions as compared to the three-buyer auctions, and these differences are significant at the 5-percent level. Table 5 presents risk-neutral Nash and double-value model tests for each of the 26 subjects who

TABLE 5—NASH AND DOUBLE-VALUE HYPOTHESIS TESTS FOR INDIVIDUAL SUBJECTS FACING ROBOT OPPONENTS (PAIRED WILCOXON SIGNED RANK TESTS, 45 OBSERVATIONS PER SUBJECT)

Ordering of the number of buyers and experience	Nash model hypothesis: Statistic too low	Nash model hypothesis: Statistic too high	Double-value model hypothesis: Statistic too low	Double-value model hypothesis: Statistic too high	Number of subjects
Inexperienced:					
363	4 (3)	3 (2)	7 (6)	0 (0)	7
636	4 (3)	4 (0)	8 (8)	0 (0)	8
All inexperienced:	8 (6)	7 (2)	15 (14)	0 (0)	15
Experienced:					
363	3 (2)	3 (2)	5 (5)	1 (0)	6
636	2 (2)	3 (2)	3 (3)	2 (1)	5
All experienced:	5 (4)	6 (4)	8 (8)	3 (1)	11
All Sessions	13 (10)	13 (6)	23 (22)	3 (1)	26

Note: The numbers in parentheses are the numbers of times the model prediction was rejected at the 5-percent level of significance (two-tailed test).

faced robot opponents (these tests employ 45 bid observations per subject). An equal number of buyers bid above and below the risk-neutral Nash prediction, and the Wilcoxon test nearly always rejects the more aggressive double-value model. Unlike the human-opponent data presented above, average bids (not shown) frequently do not exceed the risk-neutral Nash predictions.

The data shown in Table 4 indicate that 36 of the 48 human-opponent subjects (75 percent) bid more aggressively than the risk-neutral Nash prediction, and the data shown in Table 5 indicate that 13 of the 26 robot-opponent subjects (50 percent) bid more aggressively than the risk-neutral Nash prediction. These relative frequencies are significantly different ($X^2_{[1]} = 4.71$; p value $= 0.03$), using a 2×2 contingency table analysis. A 3×2 contingency table that classifies subjects as (i) bidding significantly too high, (ii) bidding significantly too low, and (iii) not bidding in a manner that is significantly different from the risk-neutral Nash prediction (all at the 5-percent level indicated in Tables 4 and 5) also indicates a significant treatment effect [$X^2_{[2]} = 5.99$; p value $= 0.05$].[20] I also tested for differences between robot-opponent and human-opponent behavior for the bids with identical value draws, pooling across subjects and again employing a random-effects model to capture systematic differences across subjects. Consistent with the results shown in Tables 4 and 5, the average robot-opponent bids are usually lower than the average human-opponent bids, and the differences are significant at the 5-percent level for all but the "experienced 636" order treatment. I therefore conclude that competing against stable, known, Nash equilibrium strategies results in less aggressive bidding than the bidding observed in the human-opponent sessions and that this result is not due to sampling (population) differences.

[20] I thank an anonymous referee for suggesting this analysis. Using a lower significance threshold for rejecting the risk-neutral Nash model results in lower X^2 statistics for the treatment differences (4.52 and 4.13 using a 10-percent or 20-percent threshold, respectively; the 10-percent critical value for this X^2 statistic is 4.61).

E. *Best Responses Against Aggressive Opponent Strategies*

The significantly more aggressive bidding in the human-opponent treatment relative to the robot-opponent treatment may result from differing best responses to opponent bids in the two environments. The robot-opponent subjects competed against Nash robots, so the bid function given by equation (11) is the best response for risk-neutral subjects. In contrast, the human-opponent subjects competed against subjects who generally bid more than the risk-neutral Nash prediction. It is easy to show that, for a risk-neutral subject bidding against $N-1$ opponents who all use the linear bid function of the form $b(v_i) = \delta v_i$, the best response is still the bid function of equation (11).[21] However, bids were constrained to be less than \$10, so linear bid functions are not possible over the entire range of values [\$0.00, \$5.83] if the slope parameter δ exceeds 1.714. For example, in the double-value model with $\delta = 2$, subjects bid the maximum \$9.99 for all value draws greater than \$4.99. I calculated best responses to the following general piecewise linear bid function:

$$(14)\quad b(v_i) = \begin{cases} \delta v_i & \text{for } v_i < 9.99/\delta \\ 9.99 & \text{for } v_i \geq 9.99/\delta. \end{cases}$$

The Nash bid function of equation (11) is generally no longer a best response if opponents employ this more aggressive bid function with $\delta > 1.714$.

[21] The objective function in this case is

$$EU(b_i) = \left[\frac{b_i}{\delta \bar{v}}\right]^{N-1} \frac{b_i}{\bar{c}} \left[v_i - \frac{b_i}{2}\right].$$

The first-order condition that characterizes the best response can be simplified to

$$v_i - b_i + (N-1)[v_i - 0.5 b_i] = 0$$

because the slope term δ cancels out, so the optimal bid function is again $b(v_i) = v_i 2N/(N+1)$.

This piecewise linear bid function provides a good approximation to the empirical bid distributions. For $N = 3$ buyers in the experienced sessions, $\hat{\delta} = 1.80$ provides the best fit ($R^2 = 0.945$); for $N = 6$ buyers in the experienced sessions, $\hat{\delta} = 1.90$ provides the best fit ($R^2 = 0.953$). The fit decreases only marginally for the double-value ($\delta = 2$) model: $R^2 = 0.937$ and $R^2 = 0.951$ for the three-buyer and six-buyer auctions, respectively. (Recall that the risk-neutral Nash bid function is $\delta = 1.5$ for three buyers and $\delta = 1.714$ for six buyers. The risk-neutral Nash model fits the empirical bid data with R^2 statistics of 0.922 and 0.947 in the three-buyer and six-buyer auctions, respectively.) The risk-neutral best response to this empirical bid function in the six-buyer auctions ($\delta = 1.90$) is calculated using numerical methods and is also piecewise linear: bid the Nash bid of $1.714 \times$value for values less than or equal to \$5.15, and bid the maximum \$9.99 for all greater values. Against the double-value bid function ($\delta = 2.0$), the value cutoff for best responses of \$9.99 falls to \$5.09. In the three-buyer auctions the risk-neutral best response to the empirical bid function ($\delta = 1.80$) is the Nash bid of $1.5 \times$value for all values; however, against the double-value bid function the best response is to bid \$9.99 for all values greater than \$5.46. Although these are not mutual best responses that constitute a Nash equilibrium, these calculations indicate that risk-neutral bidders facing aggressive opponents should bid more aggressively than the Nash model for high-value draws. Since subjects facing Nash robots do not generally bid systematically above the risk-neutral Nash prediction (Table 5), while subjects facing aggressive human opponents do (Table 4), these differing best responses may explain the significant differences in the human-opponent and robot-opponent treatments.

V. Conclusion

This paper studies a feature of the EPA auction for sulfur dioxide emission allowances that leads to potentially serious incentive problems. Sellers receive the bid price of a specific buyer, and their asking price determines their trading priority. Sellers with the lowest asking prices receive the highest bids; consequently, sellers have an incentive to submit offers that underrepresent their true costs of emission control.[22] A symmetric Nash model of the seller side of this auction (tested here for an inverted and strategically symmetric problem facing buyers) indicates that increased competition (increasing N) can make this incentive problem *worse*. Vickrey (1961) demonstrated that increasing the number of buyers in a first-price auction reduces the incentive to misrepresent true valuations. In contrast, the amount of misrepresentation in the EPA auction increases as the number of competitors on one side of the market increases, so that increased competition reduces efficiency.

The laboratory results presented here support this main qualitative prediction for buyers bidding against human opponents as well as for buyers bidding against Nash robot opponents. Increasing the number of competing buyers from three to six increased overall bids, increased winning bids, and often reduced efficiency. However, the data usually reject the quantitative predictions of the risk-neutral Nash model when buyers face other human opponents. Bids in this treatment are consistently more aggressive than the risk-neutral Nash equilibrium, although they are not as aggressive as the highly competitive "double-value" model of bidding. Unlike overbidding relative to the risk-neutral Nash equilibrium in first-price auctions (see e.g., Cox et al., 1988), the overbidding reported here is inconsistent with risk aversion. Risk aversion under the current auction rules implies bids *below* the risk-neutral Nash equilibrium. Furthermore, the asymmetry in the expected loss function suggests that overbidding should be less likely than underbidding.

However, buyers facing computerized Nash "robot" opponents did not bid systematically above the risk-neutral Nash equilibrium; an equal number of subjects bid above and below this theoretical prediction. Moreover, risk-neutral best responses to the aggressive bids in the human-opponent treatment usually involved aggressive bids above the Nash equilibrium for high-value draws. This suggests that the aggressive bids in this auction environment could be explained by a small subset of high-bidding subjects. The high bids of this subset in early periods would lead to best-response bids for risk-neutral (or possibly even risk-averse) subjects that are also above the risk-neutral Nash equilibrium. In later periods this could ultimately lead to aggressive bids for most subjects. This explanation implies that bids should increase with experience, which is strongly supported in the data. In any case, the evidence reported here indicates that the EPA auction provides little incentive for truthful value revelation and may be poorly designed for trading emission allowances.

[22] In the two-sided EPA auction, both buyers and sellers have an incentive to offer terms of trade below their true values and costs. As discussed in Cason (1993), this has an ambiguous impact on trading efficiency but produces a downward bias in transaction prices. A forthcoming paper (Cason and Charles Plott, 1996) documents this downward price bias in an experimental evaluation of the two-sided EPA auction.

REFERENCES

Cason, Timothy N. "Seller Incentive Properties of EPA's Emission Trading Auction." *Journal of Environmental Economics and Management*, September 1993, *25*(2), pp. 177–95.

Cason, Timothy N. and Plott, Charles. "EPA's New Emissions Trading Mechanism: A Laboratory Evaluation." *Journal of Environmental Economics and Management*, 1996 (forthcoming).

Cox, James C.; Roberson, Bruce and Smith, Vernon L. "Theory and Behavior of Single Object Auctions," in Vernon L. Smith, ed., *Research in experimental economics*, Vol. 2. Greenwich, CT: JAI Press, 1982a, pp. 1–43.

Cox, James C.; Smith, Vernon L. and Walker, James M. "Auction Market Theory of Heterogeneous Bidders." *Economics Letters*, 1982b, *9*(4), pp. 319–25.

______. "Theory and Behavior of Multiple Unit Discriminative Auctions." *Journal of Finance*, September 1984, *39*(4), pp. 983–1010.

______. "Theory and Individual Behavior of First-Price Auctions." *Journal of Risk and Uncertainty*, March 1988, *1*(1), pp. 61–99.

Federal Register. *Environmental Protection Agency (40 CFR) Part 73*, "Auctions, Direct Sales, and Independent Power Producers Written Guarantee Regulations; Proposed Rule." 23 May 1991, pp. 23744–59.

______. *Environmental Protection Agency (40CFR) Part 73*, "Auctions, Direct Sales, and Independent Power Producers Written Guarantee Regulations and Request for Delegation of Proposals to Administer the Auctions and Direct Sale; Rule and Notice." 17 December 1991, pp. 65592–609.

Friedman, Daniel. "Theory and Misbehavior of First-Price Auctions: Comment." *American Economic Review*, December 1992, *82*(5), pp. 1374–78.

Harris, Milton and Raviv, Artur. "Allocation Mechanisms and the Design of Auctions." *Econometrica*, November 1981, *49*(6), pp. 1477–99.

Harrison, Glenn W. "Theory and Misbehavior of First-Price Auctions." *American Economic Review*, September 1989, *79*(4), pp. 749–62.

Hsiao, Cheng. *Analysis of panel data*. Cambridge: Cambridge University Press, 1986.

Kagel, John and Levin, Dan. "The Winner's Curse and Public Information in Common Value Auctions." *American Economic Review*, December 1986, *76*(5), pp. 894–920.

______. "Independent Private Value Auctions: Bidder Behavior in First-, Second- and Third-Price Auctions with Varying Numbers of Bidders." *Economic Journal*, July 1993, *103*(419), pp. 868–79.

Montgomery, W. David. "Markets in Licenses and Efficient Pollution Control Programs." *Journal of Economic Theory*, December 1972, *5*(3), pp. 395–418.

Press, William H.; Flannery, B.; Teukolsky, S. and Vetterling, W. *Numerical recipes: The art of scientific computing*. Cambridge: Cambridge University Press, 1986.

Riley, John G. and Samuelson, William F. "Optimal Auctions." *American Economic Review*, June 1981, *71*(3), pp. 381–92.

Rustichini, Aldo; Satterthwaite, Mark A. and Williams, Steven R. "Convergence to Efficiency in a Simple Market with Incomplete Information." *Econometrica*, November 1994, *62*(6), pp. 1041–63.

Satterthwaite, Mark A. and Williams, Steven R. "The Rate of Convergence to Efficiency in the Buyer's Bid Double Auction as the Market Becomes Large." *Review of Economic Studies*, May 1989, *56*(3), pp. 477–98.

Vickrey, William. "Counterspeculation, Auctions, and Competitive Sealed Tenders." *Journal of Finance*, March 1961, *16*(1), pp. 8–37.

Walker, James M.; Smith, Vernon L. and Cox, James C. "Bidding Behavior in First Price Sealed Bid Auctions: Use of Computerized Nash Competitors." *Economics Letters*, 1987, *23*(3), pp. 239–44.