THE GEORGIA IRRIGATION REDUCTION AUCTION: EXPERIMENTS AND IMPLEMENTATION

Ronald G. Cummings, Charles A. Holt, and Susan K. Laury*

January 2002

Abstract: In April 2000, the Georgia legislature passed a law requiring that the state use an unspecified "auction-like process" to pay some farmers to suspend irrigation in declared drought years. In response, we conducted a series of experiments that tested a variety of auction procedures. This paper reports the results of the laboratory and field experiments that were used by policy makers who determined the auction procedures. The results of these experiments are compared with farmers' bidding behavior in the state-run irrigation auction conducted in March 2001.

* Cummings and Laury: Andrew Young School of Policy Studies, Georgia State University, University Plaza, Atlanta, GA 30303-3083. Holt: Economics, University of Virginia, Charlottesville, VA 22904-4182 (Holt was also a visiting scholar at Georgia State during the spring and summer of 2001). The authors gratefully acknowledge funding from the Georgia State University Experimental Economics Laboratory, the National Science Foundation (SBR-0094800), and the University of Virginia Bankard Fund for Political Economy.

THE GEORGIA IRRIGATION REDUCTION AUCTION: EXPERIMENTS AND IMPLEMENTATION

I. INTRODUCTION

In April, 2000, the Georgia Legislature enacted the Flint River Drought Protection Act.[1] The state was entering a third year of drought and Georgia=s primary water manager, the Environmental Protection Division (EPD), was concerned that the exercise of existing water use permits could reduce flows in the Flint River to levels that might cause serious harm to the Basin=s ecological systems in general, and to endangered species in particular. The Drought Protection Act was designed to reduce irrigated acreage (which accounts for more than 70% of consumptive water use in the Basin) during drought years. On March 1 of each year the Director of the EPD is required to announce whether or not the upcoming summer will be characterized by severe drought conditions. If a drought is declared, the Director is then required to determine the number of acres that must be taken out of irrigation to maintain acceptable river flows. The Act stipulates that the acreage reduction target be implemented via an Aauction-like@ process wherein farmers may offer to voluntarily[2] forego irrigation of *all*[3] lands covered by a specific water use permit. The Legislature set aside $10 million of funds derived from the multi-state Tobacco industry settlement for use in compensating farmers.

After learning of the Act's auction provisions, we immediately initiated a program of research aimed at assisting the EPD in their efforts to define the substance of the Aauction-like process@ that they might be (and ultimately were) required to implement. An appropriate design for such an auction was not immediately obvious. The necessary

---

[1] O.C.G.A. 12-5-540 through 12-5-550.

[2] The Act also provides for involuntary acreage reduction in the event that the Director=s target acreage is not achieved through the auction. The legal uncertainties surrounding this provision are formidable, however, and its implementation is problematic. For purposes of research reported in this paper, we excluded consideration of the involuntary acreage reduction provision. Indeed, in the EPD=s initial implementation of the Act with an auction conducted in March 2001 (described later in this paper) the involuntary acreage reduction provision was *not* implemented.

[3] Reflecting issues related to enforcement, farmers were not allowed to offer partial reductions in acreage covered by a specific water use permit. Offers could be made only to take all lands covered by the permit out of irrigation for the balance of the year. If, as was often the case, a farmer held more than one permit, he/she could offer to take all lands covered by one permit out of irrigation, but continue to irrigate all lands covered by permits with offers not accepted in the auction.

characteristics were somewhat unique *vis-a-vis* received auction literature in a number of ways.[4] Unlike (as contemporary examples) the SO_2 and FCC bandwidth auctions that involve a single or multiple seller(s) facing many buyers, the water use permit (hereafter, simply Apermit@) auction involves a single, *budget-constrained*,[5] *buyer* facing many sellers. Since most farmers use irrigation for one or more of three major crops, (corn, cotton, and peanuts), we concluded that sellers have relatively good common information on the value of crop yields to other sellers, subject to some uncertainty about price and weather. In contrast, differences in permit size, soil quality, and location produced variations in per-acre productivities of irrigated land, which we modeled as "private value" differences. Finally, any effort to design an auction mechanism for this application would have to reflect obvious incentives for sellers to attempt to collude.[6]

The purpose of this paper is to describe our efforts to evaluate alternative auction mechanisms to implement the Act=s requirements. The subjects were mostly students, but we made use of adults and farmers in some cases. In section II, we discuss the institutional considerations that determined the basic experimental designs and treatments. Modifications based on pilot experiments are reviewed in section III. The results presented in section IV include laboratory experiments with various scales (from 9 to 42 participants) using various pricing, bid revision, and information provision rules. Two field experiments are reported in section V, one with subjects that were primarily farmers and one that was a multi-site auction with 50 people at two locations in southwest Georgia, just two months prior to the actual auction. The bids in this trial run were collected using a web-based program that enabled the Director of the EPD to watch

---

[4] For discussions of auctions for bandwidth and pollution permits, see Cramton (2000, 2002) and the experimental results in Banks, Olson, Porter, Rassenti, and Smith (2001).

[5] To provide some degree of perspective for the EPD=s budget constraint of $10 million, in a severe drought year the EPD=s target acreage reduction could be on the order of 100,000 acres (out of a total of some 600,000 to 800,000 irrigated acres). Achievement of the target would require that the average price/acre paid in the auction not exceed $100, which is seen by many as the lower bound on rental values for irrigated land in the Basin.

[6] Nothing in the Act prohibits collusion or makes it unlawful, and we believed such collusion to be likely given the social situation at bid centers and the ease of cellular phone conversations with those at other centers. Nevertheless, an organized bidding cartel could be covered by existing antitrust laws.

the bidding from Atlanta.[7] Section VI summarizes the results of the EPD auction, which involved 194 farmers at 8 Flint River basin locations. The final section concludes.

II. Implementing the Irrigation Auction in Laboratory Experiments

Because our work was motivated by a specific policy question, the experiments that we conducted included more context and fewer controls than is normally the case. The first part of this section provides details on the institutional context for the irrigation auction. The second part outlines our laboratory experiments.

Institutional Considerations

There were many institutional details that we considered and mimicked in our experiments. First, there are many potential "sellers" (farmers) and only one "buyer" (the EPD).[8] Some farmers have more than one irrigation permit, and the value of the land is not homogenous across farmers. In fact, one farmer may hold irrigation permits for land that is used to grow different crops, and therefore has different values. The auction must be implemented and finalized quickly (within 25 days) after a drought declaration is made. Moreover, it must be able to accommodate a potentially large number of farmers, who are located over a broad area of southwest Georgia. Not all of these farmers have internet access, and some are not comfortable with computers. While the EPD will have a target number of acres that they wish to take out of irrigation, they also have a fixed budget constraint that is likely to be binding. The maximum budget is public knowledge, but the target acreage need not be. Finally, the institution should be "collusion-proof" given that many of those participating in the auction know one another and will have ample opportunity prior to (or during) the auction to discuss bidding strategies.

---

[7] The program, written by Kevin Ackaramongkolrotn at Georgia State University, was also used in the final auction. Web-based software is convenient for auctions with dispersed bidders because it can be scalable and platform independent (Lucking-Reiley, 1999).

[8] We use the terms buyer and seller to correspond to traditional auction theory and design, however nothing is being "bought" or "sold" in this auction. To be more precise (and legally accurate) farmers make offers to suspend irrigation for the remainder of the calendar year. These offers may be accepted or rejected by the EPD. Farmers, however, retain their land and their irrigation permit regardless of the outcome of the auction.

Irrigation systems in this region are not metered. Therefore it was not possible to implement a system in which a given farmer reduces irrigation to a target amount. In order to have more manageable enforcement, regulators decided that any offer to suspend irrigation required no irrigation take place on *any* of the land covered by the permit. For example, if a single irrigation permit covered three distinct fields, an offer to suspend irrigation means that no irrigation will occur on any of the three fields. However, if a farmer had more than one irrigation permit, the farmer could offer to suspend irrigation under one and still irrigate land covered by another.

An offer to suspend irrigation merely states that the farmer is willing to forego irrigating for the remainder of the calendar year in exchange for the specified (per-acre) payment. The farmer can still use the land, and in fact can still plant crops on the land. However, the land may not be irrigated, which makes it unlikely that the land will be used for agricultural purposes. Beginning January 1 of the following year the farmer is free to irrigate once again.

There were political considerations as well. The outcome of the auction should not be considered to be either "wasteful" or "unfair." There is a large variance in the quality of land (and the crops grown), both within and across regions. Many therefore perceived that a "fair" pricing system would entail paying different amounts to farmers with different values. Ultimately, these considerations led to our rejection of uniform auction formats, and the limitation of our focus to discriminative auctions. In addition, in a uniform price auction there was the potential of negative publicity if a farmer offered to suspend irrigation for a very low amount but was paid a much higher amount determined by the uniform market clearing price. While we compared uniform and discriminative pricing in our initial experiments, described below, we narrowed our focus when it became obvious to us that a uniform-price mechanism would not be seriously considered by the EPD.

Typically one considers an auction successful if the outcome is efficient. In this context, this entails not only the "right" price resulting from the auction, but also that those farmers with the lowest land values have their offers accepted. However, the EPD's primary goal was simply to take the maximum number of acres out of irrigation within the fixed budget constraint. This was particularly important given that most

thought it was unlikely that they would be able to meet their target acreage within the budget they had to work with. Efficiency was a secondary concern to them.

Finally, because an auction like this had never been conducted in this region (and is expected to happen only infrequently), it is important that the rules and procedures be clear, easily understood, and also easy to implement. In particular, it was of utmost importance that, if a farmer's offer was accepted, the farmer understand precisely how their payment would be determined.

Laboratory Implementation

In the typical laboratory experiment, neutral terminology is used for goods of an unspecified nature, in order to preclude the possibility that valuations may be biased by experiment context. Moreover, subjects are often visually isolated from one another, and are not allowed to communicate (unless this communication is a treatment variable). The complexity of the auction procedures being discussed convinced us that a fair amount of context would be useful to reduce confusion, and we did not think vague beliefs about land rental rates would affect induced values that were several orders of magnitude lower. Therefore, subjects in these experiments were told:

> In this auction, each of you will be in the position of a farmer who has three "permits" to irrigate acres of land. These irrigation permits allow you to irrigate the land and earn money on crops that you grow. We (the experimenters) are in the position of a government agency charged with controlling water use. We will use an auction-like process to buy some of these permits back from you in order to reduce the amount of water being taken from river and ground water reserves in this area.

This context-specific terminology proved to be useful when wrote up instructions for the actual auction that followed a year later.

We used induced values to determine the supply function for irrigation permits. The subject was told the (per-acre) value for each permit that was held. If the permit was not sold, the subject would definitely earn this amount of money, multiplied by the number of acres covered by the permit. If the permit was sold, the subject would not earn this money, and would instead earn the negotiated per-acre price multiplied by the number of acres covered by the permit. The certain value of a held permit is a simplification of the actual situation facing farmers, since crop price is not known with

certainty in advance. Roughly 75 percent of the acreage in this region is in corn and cotton, and this is also the acreage with the lowest profit margins. Farmers in this area are clearly price-takers in these crops (reducing acreage would have no effect on market prices), and therefore have a good idea of the costs and prices that they face. The remaining acreage is in peanuts. Because profit margins are higher, and also because of peanut quotas, we considered it very unlikely that any peanut acreage would be offered in the auction at prices close to those that were being accepted.

After the initial sequence of experiments, we were concerned that the laboratory environment inhibited the communication that would certainly occur among farmers in an actual irrigation auction. Therefore, we conducted the majority of our sessions in an open lobby area outside of the laboratory, or in a separate room, depending on the number of participants. Refreshments were provided, and subjects were encouraged to talk with one another. The auctions were conducted using a series of five-minute offer submission rounds. Subjects were told: "while you are waiting to turn in your offer submission card (or simply waiting for the end of the 5-minute period), feel free to talk with the others and to enjoy the refreshments." The experimenters were available to answer questions, but typically kept their distance (sometimes standing in another nearby room) so that they did not inhibit any conversation. During the largest experiments and in the EPD auction, the people running the auction wore red baseball caps so that they could be easily identified in the event that there were questions or problems. For the majority of experiments, subjects were students at Georgia State University.

Even before draft rules of the auction were released by the EPD, farmers in the affected area of the state were discussing the auction. We expected them to come to the auction having discussed the auction and bidding strategies amongst themselves prior to the auction. In order to better simulate this level of experience, some subjects participated in several auction experiments. About 40 percent of our subjects participated in more than one session. Therefore our auctions involved a mix of experienced and inexperienced subjects. We did, however, shift the land values (and budget constraint) by a constant between sessions in order to change the competitive price between sessions. Also, because we expected many of the farmers to know one another, we placed no restrictions on friends or family members participating together in these experiments

(which is typically avoided wherever possible in economics experiments). For example, we know of several cases where spouses, siblings, and parents participated together.

Subjects did, in fact, talk with one another during the sessions. Because communication was not a treatment *per se* (instead we were trying to parallel the naturally occurring environment) we did not monitor the conversations or keep transcripts. We did observe that subjects sometimes engaged in "small talk," and at other times talked about the auction itself. Because we wanted to ensure that any auction mechanism that we recommended would be relatively collusion-proof, we were happy that bidding strategies were discussed and attempts at collusion occurred in most sessions.

Overall, 90 subjects participated in 20 auctions held during 11 sessions in April and May 2000 (some subjects participated in as many as three auctions). Most sessions lasted for two hours (the 42-person session was scheduled to last 4 hours), and earnings ranged from $36.62 to $99.88, with an average of $63.74.

III. Pilot Experiments: Pricing Rules and Revisions

In our first pilot sessions, subjects participated in more than one auction, which allowed us to test a variety of institutions quickly before narrowing our focus to one or two sets of rules.

The design and land values used in the last two pilot sessions most closely resembled those conducted subsequently. Therefore, we will report only the results from these sessions. Subjects in both participated in two sealed-offer auctions: a discriminative (own-price) and uniform-price auction. They also participated in a final (uniform-price) auction in which they were allowed to revise their offers after preliminary results were announced. In one session the uniform price was set equal to the lowest rejected offer price, and in the other session it was set equal to the highest accepted offer price. We were interested in how the average price paid, number of acres obtained in the auction, and efficiency were affected by the choice of institution and the opportunity to revise offers based on market information. In order to minimize order and information effects we first conducted the one-shot (no revision) auctions without providing any feedback on the results. Therefore when the iterative auctions were conducted, subjects had

previously submitted offers, but had not observed the outcome from either one-shot auction.

A uniform price auction is typically preferred because it is incentive compatible to bid one's own value (at least when one's bid cannot determine the market-clearing uniform price). However, because subjects in our experiments (like farmers who were to be participants in the auction) had multiple permits, neither pricing rule we tested is theoretically incentive compatible (Smith, Williams, Bratton, and Vannoni, 1982; Ausubel and Cramton, 1998). To see this, think about a farmer making offers for two permits. If one offer is accepted, it is possible that his offer on the second (rejected) permit could determine the market-clearing price. In this (multi-unit) environment bidding one's value is incentive compatible only if the price is determined by the lowest rejected offer that is not one's own.[9] We did not consider such a pricing rule, however, because of the complexity of explaining and implementing it. Moreover, it wasn't clear to us that farmers would understand the incentive to bid value in such a complicated environment.

In each auction subjects knew their own land values, the range of values, the number of participants, and the fixed budget. They were told that we had a target number of acres that we wanted to take out of irrigation, and that we would accept as many offers as possible until we either reached this target or expended our budget. Subjects were not told the target number of acres. Moreover, in the auctions with revisions, subjects did not know how many revision rounds would be conducted.

In these auctions we did not observe a consistent difference between the uniform and discriminative auctions. In one session, the uniform price was below the average price paid in the discriminative auction. However, two accepted offers were substantially higher than others in the discriminative auction, pulling up the average. In the second session the average prices were identical. More evidence comparing the uniform and discriminative auctions will be presented below. Similarly, there is little difference in the uniform price auctions between a price based on the highest accepted offer and the lowest rejected offer.

---

[9] List and Lucking-Reiley (2000) document such withholding behavior in field experiments.

The basic procedures of the auction with revisions were the same as those for the uniform price auction. After all offers had been submitted in writing, they were ranked from low to high. The lowest priced offers were then "provisionally" accepted. After the provisional winners were announced, all subjects (regardless of the status of their offer) were given the opportunity to turn in a revised offer. If no new offer was turned in, the previous offer stood. The new offers were then ranked, and new provisional winners were announced. This process continued until either no one wished to submit a revised offer or the experimenters chose to end the auction. In this case, the provisional acceptances from the most recently completed round became final acceptances. Subjects did not know in advance which would be the final offer round.

We placed no restrictions on the revisions. Therefore a subject who initially submitted a high offer price could lower their offer (provided another revision round was held). Similarly, a subject who submitted a very low offer could increase it, even if the offer was provisionally accepted. Of course, doing so involved the risk that the subject would be excluded from the market at the new offer price.

A key issue here is how to announce the provisional winners. In particular we considered whether we should simply announce which offers were accepted (identified in an anonymous manner, for example by permit ID number) or announce the cut-off offer price that determined which offers were accepted. For simplicity we chose to announce the cutoff price ("All offers at or below $1.20 were provisionally accepted"), however, this is a treatment that we consider below.

Across revision rounds, the average accepted offer price generally declined (from $1.11 in the first round to a low of $1.04 in round 3) in one session, and was flat in the second session (see Figure 1). In both of these auctions the uniform clearing price was *below* the competitive prediction of $1.10, obtained by intersecting the supply function (locus of ranked opportunity costs) with a demand function that is the locus of points where the price multiplied by the number of acres exactly matches the budget constraint. This sub-competitive result reflects the fact that some subjects were making offers below value, which suggests that subjects were almost certainly confused about how their earnings were determined. We addressed this in follow-up experiments, described below.

<Figure 1 about here>

The most important result to come out of these initial experiments is that inefficiency decreases dramatically when subjects are given preliminary results and allowed to revise offers. Our measure of inefficiency is the amount by which the opportunity cost (the value of a permit that is kept) of the accepted offers exceeds the minimum opportunity cost of the number of acres accepted in the auction. If the lowest-valued permits are obtained in the auction (regardless of the price paid for them) this measure will equal zero. When higher-opportunity cost permits are obtained instead of lower opportunity cost units, this measure will increase. Inefficiency is shown on the right axis of Figure 1, and peaks in the first or second revision round. By the final (fourth) revision round the "right" (minimum opportunity cost) offers are typically accepted.

As noted earlier, efficiency was not one of the EPD's primary goals. A one-shot sealed offer auction would be much easier to implement. Farmers could mail in their offers, which would then be sorted by policy-makers. We need a strong justification for the additional time and expense of conducting an auction with revisions (which by necessity involves taking workers into the field). One reason is simply to insure against a bad outcome. Given that this type of auction had never been conducted before, policy-makers would like some assurance that prices would be at a “reasonable” level. If, for example, a simple sealed offer auction is conducted and the submitted offers are extremely high, very few offers would be accepted. Allowing for revisions gives farmers a chance to think about the situation and to respond to policy-makers’ decisions and the bidding behavior of others. Allowing for these revisions could minimize the chance that farmers will come out of the auction wishing that they could do something differently. Given the potential political repercussions of a poor outcome, or unhappy farmers, this is a big advantage of implementing an auction with revisions. Moreover, holding other factors constant, an institution that results in more efficient outcomes (and therefore, presumably, participants who are more satisfied with the outcome) is preferable. This is a strong argument in favor using an auction mechanism that includes the opportunity to revise offers.

IV. LABORATORY EXPERIMENTS

This section describes the treatments that we tested in the lab, and the results from auctions using each of these treatments. Among the factors considered were: the tie-breaking rules, uniform versus discriminative pricing, number of participants and the information provided about cutoff offers.

Training Subjects to Understand the Institution

After the pilot sessions, all other laboratory sessions involved a single auction. After observing that some subjects consistently bid below value in our pilot experiments, we were concerned that some participants might be confused about how their earnings were determined in the auction. There were two potential sources of confusion: how much they earned if they retained a permit, and how their earnings were determined in the event that an offer was accepted. In a typical experiment, the auction is conducted over a series of rounds, and earnings (for sold and unsold units) are reported at the end of each round. This was not the case in our experiment. If a subject was confused about how earnings were calculated, they received no information during the experiment that would eliminate this confusion.

Because of this we used extensive instructions (contained in the appendix) to explain the procedures and how earnings were calculated. Participants were asked to calculate their earnings for each permit that was not sold in the auction. In addition, the experimenter publicly worked through examples of how earnings would be calculated if offers on no permits or some permits were accepted. These examples used prices that were quite different than any potential prices in the market. Finally, participants worked through two practice auctions using real goods (for example, pens or post-it notes). In the first practice auction, subjects were endowed with one unit of a good (for example, a single pad of post-it notes). After the fixed budget was announced (typically several dollars), subjects submitted the price at which they would be willing to sell the good back to the experimenter. Offers were publicly recorded and ranked on a transparency and the lowest-priced offers were accepted until the budget had been expended. Those participants whose offers were accepted were paid (either the uniform or discriminative price, depending on the treatment) and the item was taken from them. It was emphasized

that those participants whose offers were accepted received the money but not the value to them of the good. The others received no money but were still able to use the good. The second auction was identical to the first, except that subjects were given two possibly heterogeneous units of the good (for example, a blue pen and a black pen) to more closely correspond to the multi-unit auction design. While this did not keep all subjects from bidding below value in subsequent auctions, very few cases of this were observed.

Tie-Breaking Rules

Given a fixed budget constraint, the possibility exists that a tie could occur at the highest accepted price. For example, suppose the budget was $10 and the ranked offers in a discriminative auction were: $0.50, $1.00, $2.00, $3.00, $3.00, $3.00, $5.00, and $10.50. In this case, we could accept all offers below $3.00, but only 2 of the 3 offers at $3.00.

In one (uniform-price) session, a subject asked how we would choose which offers to accept if there was a tie such as the one described above. We responded that we would accept all offers at this price, "even if we have to go a little above our budget." Figure 2 shows the results from this session (and a paired session, described below). The induced supply and demand arrays are shown on the left side of the graph. The supply curve is simply the permit values, ranked from low to high. The demand side of the market comes from the experimenters, and is determined by the fixed budget constraint. For example, in this market the budget constraint was $160. Therefore we could afford to buy one acre at a price of $160 per acre, 4 acres at a price of $40 per acre, or 160 acres at a price of $1 per acre. The demand curve traces out the locus of these points where the price multiplied by the number of acres exactly matches the budget constraint. In this session, the uniform competitive price was $1.10 (prices on the graph are shown in pennies). We define this price as the one at which the number of acres that would be offered in the market if all bid value (144 acres) is just what we could afford to purchase at a uniform price within our fixed ($160) budget.[10]

[10] If we assume that all will bid 1-cent above value in order to avoid indifference, the price would be $1.11, and 144 acres could still be retired while staying within the fixed budget.

The right side of Figure 2 shows the time series (across revision rounds) of average accepted prices obtained in this market (shown in the "inclusive tie-breaking rule" line).[11] In the first offer round, if we had observed our $160 budget constraint, we would have accepted two of these offers (plus six offers at lower prices), for a total of 112 acres at a cost of $134.40. However, because of our announced procedure of accepting all offers at the tied level, we provisionally accepted offers for 160 acres at a total cost of $192. In the second round the highest accepted offer price fell to $1.15, and there was once again a tie. As additional revision rounds continued, almost all permits with a value below $1.15 were submitted at this level – even those with very low values in this uniform price auction. By round 6, there were 18 accepted offers. All but five of these were submitted at $1.15. We retired a total of 288 acres (double the competitive level), and spent $331.20: over double our budget. Extrapolating to the actual problem faced by the EPD, this translates into spending over $20 million when the available budget is $10 million.

<Figure 2 about here>

Later on the same day we conducted a second session (with a different group of participants). This session was identical (number of participants, parameter values, and procedures), except for an announcement that, in the event of a tie, we would randomly choose among offers at the tied level in order to stay within our budget constraint. The initial offers were quite similar to the first session (and are shown as the "random tie-breaking rule" line in Figure 2). In round 1, the highest accepted offer was at $1.15. There was a tie, and we accepted 2 of 3 offers at this level. In each subsequent round the maximum accepted offer fell. In the end, the price as 2-cents below the competitive level, and we were able to retire 144 acres.[12]

---

[11] Because the highest accepted offer is the one (uniform) price paid for all accepted offers, this is also the average price paid. For consistency with subsequent figures that show data from discriminative auctions, we label this as the average price.

[12] The price was below the competitive level of $1.10 because one subject offered a permit at a price below value, and in fact the final price was below the value for this permit. While this sometimes occurred in these sessions, this was not typical of bidding behavior.

Uniform versus Discriminative Pricing

Next, we tested the effect of the pricing rule used on offers and average prices paid in these auctions. Figure 3 shows the results of several auctions run with the same budget, value arrays, and random tie-breaking rule. We observed a clear tendency for the range of offers to lie above values in both uniform and discriminative auctions, especially in early offer rounds (see the top panel of Figure 3). There was little difference in the median offer-to-value ratio among accepted offers in the two types of auctions. In each of these auctions, the offer-to-value ratio *increased* across revision rounds. In the uniform-price auction, the median ratio increased from 1.02 in Round 1 to 1.04 in Round 6. Combining the two discriminative auctions, the median ratio increased slightly: from 1.02 in Round 1 to 1.03 in Round 5.

<Figure 3 about here>

In each of these experiments, the maximum accepted price decreased across revision rounds. In the uniform-price auction, this is maximum accepted price is the one price that is paid for all accepted offers. Therefore, we can say that the average (uniform) price fell across revision rounds in the uniform-price auction. In contrast, the average price in the discriminative auction typically *increased* over revision rounds. Many subjects whose offers were provisionally accepted in early rounds raised their offer price, resulting in this increase in the average price paid as more offer submission rounds were held. By the final rounds of the discriminative auction, most accepted offers were at or near the market-clearing price, effectively removing any advantage of a discriminative pricing rule (see the bottom panel of Figure 3). So, while average prices are initially lower in the discriminative auction, this difference tends to diminish or disappear as bidders are allowed to revise their offers. This is shown in Figure 4, which displays the average prices paid for these auctions.

<Figure 4 about here>

At this point in our research it became clear to us that the uniform price auction wouldn't be seriously considered by the EPD. Therefore all remaining experiments used the discriminative auction.

Scale of the Experiment

All of our initial experiments were conducted in groups of nine participants. However, we wanted to ensure that the procedures and results were robust to changes in the scale of the experiment. Using the discriminative auction with revisions (randomly choosing among tied offers), we conducted one session with 20 participants and another session with 42 participants. None of our key results were affected by this increase in the number of participants. The highest accepted offer declined across revision rounds while the average price paid generally increased. However, average prices remained near competitive levels despite attempts to collude. Figure 5 shows the average price paid across offer revision rounds in the 20-person and 42-person discriminative auctions. The competitive price in the 20-person auction was $1.26, and 40-cents higher in the 42-person auction. This increase in values was done both to disguise the competitive price for those who had participated in a previous auction, and also to increase payoffs for this auction (which lasted almost four hours). For comparability, 40 cents was subtracted from the average price paid in the 42-person auction before constructing the graph.

<Figure 5 about here>

Attempts at collusion in these experiments were quite explicit. Some subjects stood up to address the group, encouraging all to submit high offers. People worked together in groups, and at times a single person would turn in offer submission cards for all of those in the group.

Information About Cutoff Offers

In retrospect, the increase in average accepted price over revision rounds in our discriminative auctions made sense. Those who initially submitted very low offer prices observed the highest accepted offer, and therefore they had clear information on the price others were receiving in the auction. The incentive for low-valued participants to increase their offer was clear, and they responded to this incentive. A typical pattern of behavior is shown in Figure 6, which shows the offer price submitted for 3 participants in the 42-person auction experiment. After receiving a provisional acceptance, bidders typically raise their offer price. Sometimes this increase is gradual; others increase their offer by a large amount. After being excluded in a subsequent revision round the subject

then tends to decrease the offer again in order to get back into the market. Given that the maximum accepted offer typically falls across revision rounds, some subjects never again receive a provisional acceptance (as was the case for Subject 1).

<Figure 6 about here>

Because the announcement of the highest accepted offer had this effect, we conducted additional (small-group) sessions in which we announced accepted offers (identified by permit ID number), but did not announce the highest accepted offer price. Figure 7 shows the average price paid in one of these sessions. As anticipated, we observed an initial decline in the average accepted offer price. However, we were surprised to see this average price begin to increase again (very gradually) after several offer revision rounds. Looking at the individual data, however, helps to explain this. Subjects were quite sophisticated in how they used information. Recall that each subject had three permits with heterogeneous values. Typically, a participant offered each permit at a different price. Therefore, if two offers were accepted, this gave an upper-bound on the amount by which the lower offer could be relatively safely revised upwards. If only one offer was accepted, the participant often experimented with the offer price on one or both permits to ascertain the highest accepted price. Still, given the initial decline in offer prices and the fact that some risk averse subjects did not revise provisionally accepted offers, we concluded that this may be a more effective way to release information about which offers were accepted.[13]

<Figure 7 about here>

V. FIELD EXPERIMENTS

Two experiments were conducted in southwest Georgia, using somewhat different subject pools. The first utilized 22 adult subjects in Albany, Georgia. Most of these participants were farmers in the affected area. In order to simplify procedures, and expedite instructions, subjects were each given two "vouchers." Each voucher had a redemption value printed on the face of it. If the voucher was retained the subject received this redemption value in cash. If it was sold, the participant earned his or her

[13] See Goeree, Holt, and Palfrey (2002) for references and new evidence relating to risk aversion in private value auctions.

offer price. This is equivalent to a permit that covers a single acre (so that the per-acre price received is identical to the redemption value for the voucher).

These subjects first participated in a sealed bid (no revision) discriminative auction. After turning in an offer, and before any results were announced, they next participated in a discriminative auction with revisions. No information on the highest accepted offer was released: only those permits whose offers were provisionally accepted. A fixed budget constraint was not used in this session: the lowest 15 offers were accepted without consideration of the amount of money it took to purchase these vouchers.[14] The distribution of voucher values (which was approximately uniform in this auction, as shown on the left side of Figure 8) was also different from previous auctions.

<Figure 8 about here>

Data from these auctions are shown on the right side of Figure 8. The average accepted offer was higher in the first round of the auction with revisions than in the one-shot sealed offer auction ($21.05 compared with $19.80). Moreover, the opportunity cost of obtaining 15 vouchers was only 7.5 percent higher than the minimum opportunity cost in the one-shot compared with 12.5 percent in the first offer submission round of the auction with revisions. However, after round 1 the average accepted price was lower in each revision round than in the one-shot auction. With one exception (round 6) the opportunity cost of obtaining these vouchers was lower in the auction with revisions. The average accepted price dropped dramatically between the first and second offer rounds (from $21.05 to $19.25). The average accepted offer was at its minimum in Round 3 ($18.71), but only increased slightly after this, remaining fairly steady just under $19.00 through the remaining five rounds. This pattern is quite like that observed in our other auctions in which the maximum accepted offer was not announced (see Figure 7, above).

By late January it appeared likely that a drought would, in fact, be declared. We therefore conducted a multi-site experiment that field-tested the auction preparations. We had several concerns that were addressed in this field experiment. First, we wanted to

[14] When the EPD released the initial auction rules for public comment it proposed either a sealed offer auction, or a sealed offer with revisions. This session was held after the release of these draft rules to educate farmers on the proposed rules, and also to obtain data that compared the two institutions. The farmers were paid for participation.

determine what facilities would be needed for the EPD auction. In addition, this was our first large-scale test of the auction software. This experiment allowed us to ensure that the software could handle data entry at multiple locations, with a variety of computer systems, and expeditiously transfer information about bids and provisionally accepted offers between the auction sites in Southwest Georgia and the central processing location in Atlanta.

This field experiment was conducted at two sites in southwest Georgia, all located within the Flint basin.[15] Most of the participants were high school and college students. However, some farmers (interested in observing a live demonstration of the auction procedures) also participated in the auction. A total of 50 subjects participated in the field test, with bid collection and processing done via a web-based program that enabled the EPD officials in Atlanta to follow the bidding.

The parameters and procedures were identical to the Albany field experiment (except that a budget constraint was enforced). All subjects were given two vouchers at the start of the auction. Each voucher represented a single acre, and the values were approximately uniformly distributed from $15.00 to $22.50. We utilized a target of 55 vouchers, and a budget of $975 to purchase these vouchers. As in our laboratory experiments, the budget was common information among all participants, however the voucher target was not announced. At the competitive (uniform) equilibrium, 50 vouchers could be purchased at a price between $18.50 (the value of the last four vouchers) and $19. This is just short of the target number of vouchers (55). Of course, because we were using a discriminative auction it is possible that more offers could be accepted if the average accepted offer was less than $19. After all offers were submitted at both sites, they were combined and ranked in order from lowest to highest offer price. Starting with the lowest prices, offers were provisionally accepted until either 55 vouchers were obtained or the cost of obtaining another voucher put the total cost above $975. In the case of a tie at the cutoff value, offers at this level were randomly chosen for provisional acceptance. Provisional winners were posted, using the permit ID number associated with accepted offers. No information about the cutoff value was announced.

---

[15] As described below, we intended to conduct this experiment at three locations.

<Figure 9 about here>

On average, offers in this treatment started very low. In the first round, 55 offers were accepted at an average price of $17.58. However, there appeared to be some confusion among the subjects: almost 20 percent of all offers were below value (see Figure 9). Over time, however, the subjects appeared to learn about the incentives, quite possibly through conversations with other auction participants. In the final three rounds 50 offers were accepted – the competitive prediction – at an average price between $19.30 and $19.35 in each of these three rounds. The opportunity cost of these 50 vouchers was 7 percent above the minimum possible to obtain 50 vouchers. Individual behavior was very similar to that observed in our lab experiments. Across revision rounds, those subjects who submitted initial low offers increased them (see the left side of Figure 9), while those who submitted high offers reduced them (right side of Figure 9). By the final offer submission round, the distribution of offers was close to uniform.

There were some unexpected procedural difficulties during this trial auction. We intended to use three sites for this auction, but officials at one location forgot to send someone to unlock the building. We discovered problems with computer "firewalls" that inhibited communication between sites (specifically with a chat-room that was set up for site-supervisors to communicate with one another). Moreover, we expected to use cellular phones to communicate between auction sites and the central processing location in Atlanta; however, cellular coverage was unreliable at all of the remote sites. Finally, we encountered trouble with the software that was used to print receipts (showing the final auction outcomes and payments to be received). This trial auction provided us with valuable guidance about remaining preparations that needed to be completed. However, it also convinced us that a multi-site auction with a diverse collection of computer facilities and a large number of bidders was feasible.

VI. THE FLINT RIVER DROUGHT PROTECTION ACT AUCTION

The EPD's Flint River Drought Protection Act Auction was conducted on Saturday, March 17, 2001 at eight sites in the Flint River Basin.[16] Two weeks prior to the

[16] There was one site supervisor and two EPD representatives at each of the eight auction sites. In addition we hired a total of 58 people to work at the eight locations. These workers collected bids from farmers,

auction, all eligible participants were sent instructions that detailed the auction procedures and directed them to the eight sites.[17] On the day before the auction, eligible permit-holders could register to participate at any one of the auction locations. A total of 576 permits (covering 98,170 acres) were certified as eligible for the auction. Of these, about two-thirds were registered to participate in the auction. A total of 194 farmers registered to make offers for 347 permits, totaling 61,806 acres. The acreage associated with these permits ranged from 4 to 1442 acres. Although we have some information about the crops (and associated prices) in this part of the state, we do not observe the values that the farmers associate with each irrigation permit. Instead, we can only observe the offers that they make on each permit. In all rounds, the per-acre offer prices ranged from $0.01 to $8,000. Arguably, the offers at these extremes weren't serious offers. In fact, the person who made the 1-cent per acre offer (for a permit that covered 20 acres) stated that he was doing so as a protest.[18] The $8,000 per-acre offer was for a permit associated with a four-acre tract of land. About 85 percent of the acreage in the auction was offered at prices from $100 through $500.

<Figure 10 about here>

Figure 10 shows the array of offers, for those offers that range up to $500 (this comprised just over 90 percent of all offers). Figure 11 shows a close-up view of these offers; the three panes divide offers into low (less than $130), medium ($130 - $210), and high ($220 - $500) offers.[19] Over the revision rounds offers did, typically decline (though there were some small increases at the low-end of the offer arrays, especially in the third and fourth rounds). Although the maximum accepted offer was not announced (only the permit ID numbers of those offers that were provisionally accepted), some

---

entered bids on the computer, and worked with farmers as they verified that offers were entered correctly. We gratefully acknowledge the help of Maribeth Coller, who helped lead auction preparations, trained these workers, and supervised an auction site. We also thank the other site supervisors: Lisa Anderson, Paul Ferraro, Ann Gillette, Laura Taylor, and Mark Van Boening (Holt and Laury supervised the remaining two auction locations).

[17] These instructions are available on the web at http://www.gsu.edu/~ecoskl/research.htm.

[18] This person never cashed the resulting 20-cent check that he received. Instead, he had it framed and has been pictured with it in several news stories, while criticizing the auction.

[19] Very high offers (those greater than $500) changed very little across revision rounds.

farmers communicated both within and between the eight auction locations. For example, at the Webster County site (where approximately 20 percent of the permits and acreage were offered) some participants encouraged others to submit at an offer price of $200 per acre. Moreover, they asked one another whether their offers had been accepted, and at what price. Several had cellular phones and called people at other auction locations in order to obtain the same information. Therefore, we can conclude that they had at least some information (though not perfect) about the range of accepted offers. In the fourth round, there were 55 offers submitted at $125 per acre, but only 42 of these were accepted. Several participants apparently forgot the random tie-breaking rule that was described in the mailed instructions; participants at several sites questioned why some offers were accepted at this level, but not others. The tie-breaking rule was publicly explained at these auction locations. Therefore, it is quite likely that many participants knew that $125 was the maximum accepted price in the fourth round and that not all offers at this level were accepted. In fact, 14 of these offers were lowered in the next (and final) round. Surprisingly, these 14 offers were about evenly divided between offers that had been provisionally accepted (8 offers) and rejected (6 offers) in Round 4.

<Figure 11 about here>

During the auction, the EPD director made all decisions regarding the conduct of the auction. In particular, he chose the rule used to determine which offers were accepted, and how many offer submission rounds to hold. Unlike the laboratory experiments that we conducted (and over our strenuous objections), he did not use a fixed budget, acreage target, average price, or maximum accepted price during the auction. This is shown in Table 1, which lists the maximum and average accepted price, the cumulative number of acres, and the cumulative cost of all provisionally accepted offers during each round of the EPD auction. There was little change in the maximum accepted offer price in the first four offer rounds. However, the EPD director increased the average accepted offer price from $105 in the first round to $113 in Round 4, increasing the total number of acres that were provisionally accepted.

<Table 1 about here>

We were concerned that the round-by-round increase in the average accepted offer price would reduce competitive pressures in the market. In fact, the number of

acres covered under provisionally accepted offers *decreased* between rounds 3 and 4, while the average price was essentially held constant (at $112.36 in round 3, and $113 in round 4). Anticipating the end of the auction, we encouraged the EPD Director to remain firm in the maximum accepted offer price in round 4, hoping this would encourage farmers to reduce their offers in the subsequent offer submission round. In fact (as noted above) this happened. The fifth line of Table 1 shows the result that would have been obtained had the maximum average price of $113 been enforced in round 5. Over 3,500 more acres would have been taken out of irrigation than in round 4, and the maximum accepted price would have been unchanged at $125 per acre. In reality, the EPD director chose to accept all offers through $200 an acre in the fifth (final) offer submission round (see the last line of Table 1). In the end, a total of 33,006 acres were taken out of irrigation at a total cost of almost $4.5 million (an average price of. $135.70 per acre).

Because a consistent rule was not used during the auction it is difficult to directly compare the outcome of the auction across revision rounds (or compare it to the results from our experiments). However, by fixing a rule and observing what the outcome using this rule *would have been* in each round we can approximate this analysis. Table 2 shows this for several different rules that might be used in an auction of this type – target acreage, average accepted price constraint, and a fixed budget constraint.

<Table 2 about here>

The top section of the table compares results that would have been attained if an acreage target had been used. For example, in the first line the target number of acres is 8,000 (just over the actual number obtained in the first round of the auction). If this had been the EPD director's goal, he could have attained 8,000 acres at an average price of $107.67 per acre in round 1, and $99.18 per acre in round 5. The number in bold (in this case, $98.90 in round 3) shows the best outcome that would have been attained in any offer submission round. The next two lines show the same comparison for higher acreage targets. The middle section of the table assumes an average price constraint was used, while the last section of the table assumes a fixed budget constraint was in effect. For all three sets of comparisons, constraints were chosen that were consistent with the actual targets used in the first, middle, and late rounds of the auction. For example, in

Round 1 the average price of provisionally accepted offers was $105, in Rounds 3 and 4 the average accepted offer was close to $113, and $136 in the final offer round.

With one exception (8,000 target acreage) the best outcome that would have been attained using any of the three rules would have been achieved in Round 5. This demonstrates the benefit of allowing farmers to revise their offers. Even though the distribution of offers did not change substantially (Figures 10 and 11), those changes that did occur generally allowed the EPD to obtain a greater number of acres at a lower price. Holding the average accepted price essentially constant between Rounds 3 and 4 was helpful in lowering Round 5 offer prices. By most measures, the Round 4 outcome was worse than that observed in Round 3. However, as we note above, offer prices decreased in Round 5, which led to the improved final-round outcome (relative to any previous offer submission round).

VII. CONCLUSIONS

Auctions are commonly used for perishable commodities like fish and flowers, and they are also used in public settings where fair access is important. Auctions can be desirable relative to administrative proceedings, since the bids convey important private information about value. The use of an auction for irrigation reduction in Georgia was attractive for these reasons, and in particular because of the narrow time window between the state-mandated drought declaration date (March 1) and the March 25 deadline to finalize the auction outcome. The auction also let farmers' bids reveal (at least indirectly) their willingness to forego irrigation on designated tracts of land for the current growing season. This avoided the anguish, inefficiency, and administrative problems of involuntary usage shutdowns using non-economic (geographic and precedence) criteria.

There are many ways to set up an "auction-like process" called for in the legislation, and we used laboratory experiments to sharpen our thinking on a number of issues: the pricing rule (uniform or discriminatory), the closing rule (with or without bid revision rounds), and how provisional results would be reported after each bid revision round. The auctions being envisioned were relatively complex environments for the bidders (student subjects and farmers), and we devoted considerable effort to coming up with a procedures and instructions that were relatively easy to understand. The laboratory

experiments enabled us to make recommendations about rules (on tie breaking and information provision) that augmented competition, even in laboratory situations where socializing and collusion were facilitated.

After attending our 42-person laboratory experiment, the EPD implemented our recommended auction institution: an iterative discriminative auction. A multi-site field experiment provided a glimpse of how farmers would behave, and of what was needed to scale up the web-based procedures to handle hundreds of bidders at eight locations in Southwest Georgia. The EPD auction held on March 17, 2001 resembled the laboratory and field experiments in some (but not all) respects, and it was considered a success.

REFERENCES

Ausubel, Lawrence M. and Peter Cramton (1998) "Demand Reduction and Inefficiency in Multi-Unit Auctions," Discussion Paper, University of Maryland.

Banks, Jeffrey, Mark Olson, David Porter, Steven Rassenti, and Vernon L. Smith (2001) "Theory, Experiment and the Federal Communications Commission Auctions," *Journal of Economic Behavior and Organization*, forthcoming.

Cramton, Peter (2000) "A Review of *Markets for Clean Air: the US Acid Rain Program*," *Journal of Economic Literature*, 38 (September) 627-633.

Cramton, Peter (2002) "Spectrum Auctions," forthcoming in Martin Cave, Sumit Majumdar, and Ingo Vogelsang, eds., *Handbook of Telecommunications Economics*, Amsterdam: Elsevier Science B.V.

Goeree, Jacob K., Charles A. Holt, and Thomas Palfrey (2002) "Quantal Response Equilibrium and Overbidding in Private Value Auctions," *Journal of Economic Theory*, forthcoming.

List, John and David Lucking-Reiley (2000) "Demand Reduction in Multiunit Auctions: Evidence from a Sportscard Field Experiment" *American Economic Review*, vol. 90, no. 4 (September), pp. 961-972.

Lucking-Reiley, David (1999) "Using Field Experiments to Test Equivalence Between Auction Formats: Magic on the Internet," *American Economic Review*, vol. 89, no. 5, 1063-1080.

Smith, Vernon L., Arlington W. Williams, W. Kenneth Bratton, and Michael G. Vannoni (1982) "Competitive Market Institutions: Double Auctions vs. Sealed Bid-Offer Auctions," *American Economic Review*, vol. 72 no. 1 (March) 58-77.

Table 1. Results from the March 17, 2001 Irrigation Auction

Offer Round	Maximum Price	Cumulative Acres	Cumulative Cost	Average Price
1	$130	7,311	$766,771	$104.88
2	$127	12,755	$1,401,843	$109.91
3	$127	17,061	$1,917,036	$112.36
4	$125	15,854	$1,791,449	$113.00
5	$125	19,406	$2,192,789	$113.00
	$200	33,006	$4,478,842	$135.70

Table 2. Results from the March 17, 2001 Irrigation Auction

Results Under Alternative Cutoff Rules

Key: Bold figure is the best outcome from any offer round

Acreage Target	Average Cost of Acquiring Target Number of Acres				
	Round 1	Round 2	Round 3	Round 4	Round 5
8,000	107.67	101.90	**98.90**	102.08	99.18
16,000	132.53	115.14	111.61	113.14	**110.59**
32,000	177.78	152.43	140.35	137.57	**113.94**

Average Price Constraint	Number of Acres That Would Be Obtained				
	Round 1	Round 2	Round 3	Round 4	Round 5
$105	7,311	9,604	10,677	9,436	**11,386**
$113	8,977	14,886	17,061	15,740	**19,406**
$136	17,110	25,130	29,912	31,081	**33,006**

Fixed Budget	Average Cost of Acquiring Acres Obtained Within Fixed Budget				
	Round 1	Round 2	Round 3	Round 4	Round 5
$1 million	111.31	104.64	102.81	104.91	**102.42**
$2 million	130.01	116.79	112.36	114.01	**111.89**
$5 million	170.26	152.43	145.22	143.58	**141.74**

Figure 1. Comparing Uniform, Discriminative, and Uniform Auction with Revisions

(average price of accepted offers and inefficiency in each offer round)

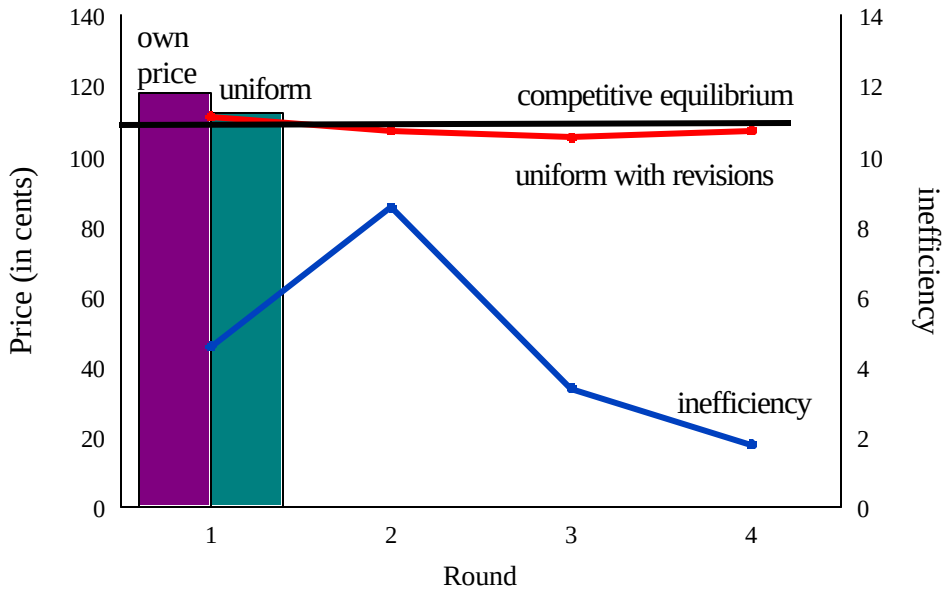

comp. price = 110 or 111

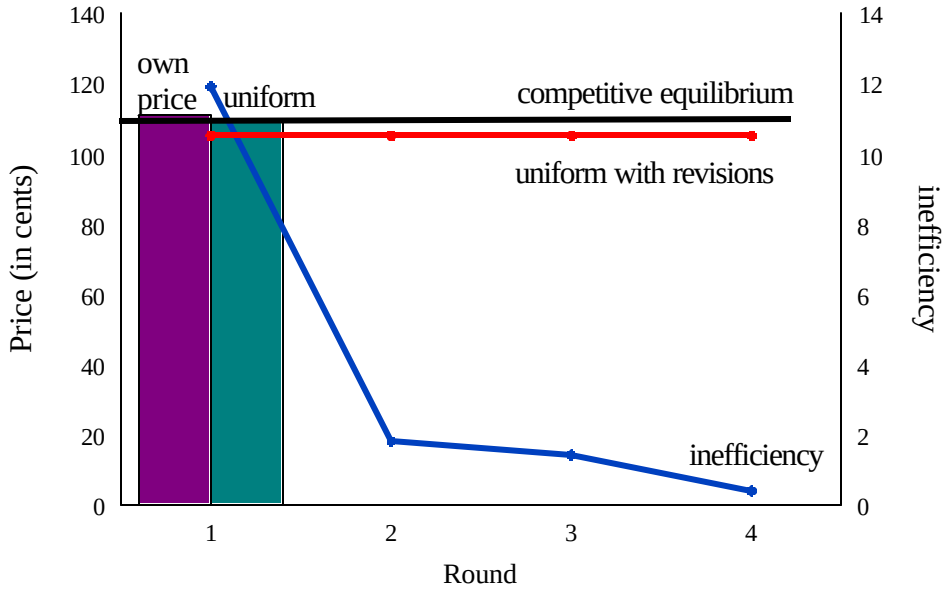

comp. price = 110 or 111

Figure 2. The Effect of Tie-Breaking Rules in a Uniform-Price Auction
Average Price Paid in Two Auctions

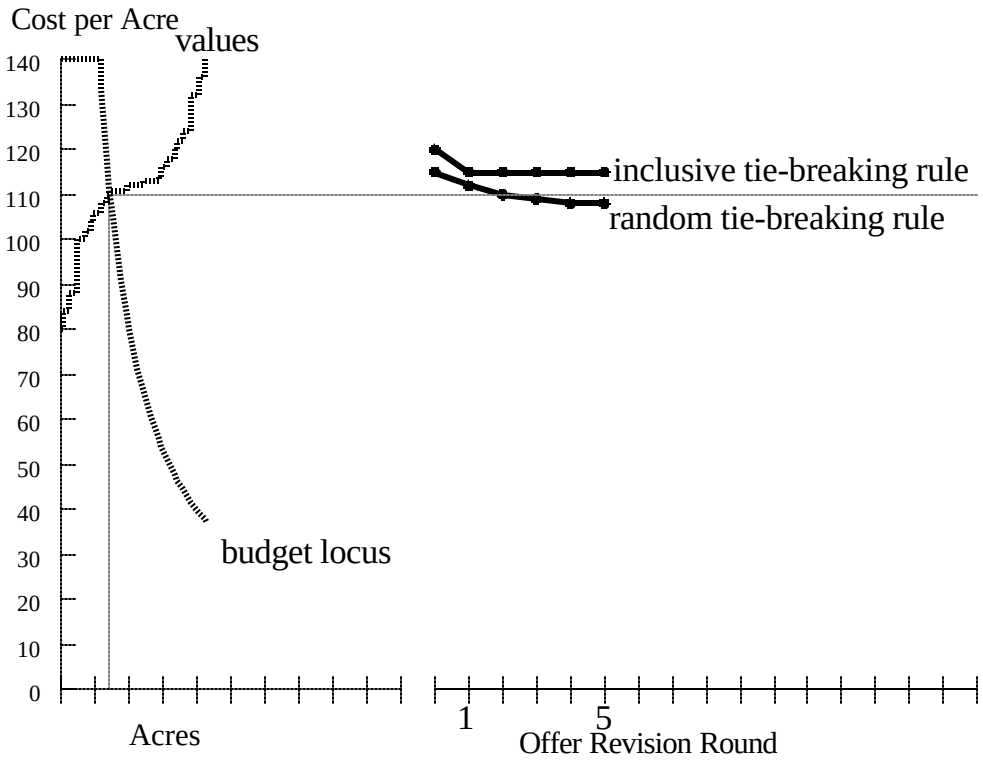

Figure 3. Distribution of Offers: Uniform and Discriminative Auctions
key: uniform (crosses); discriminative (dots)

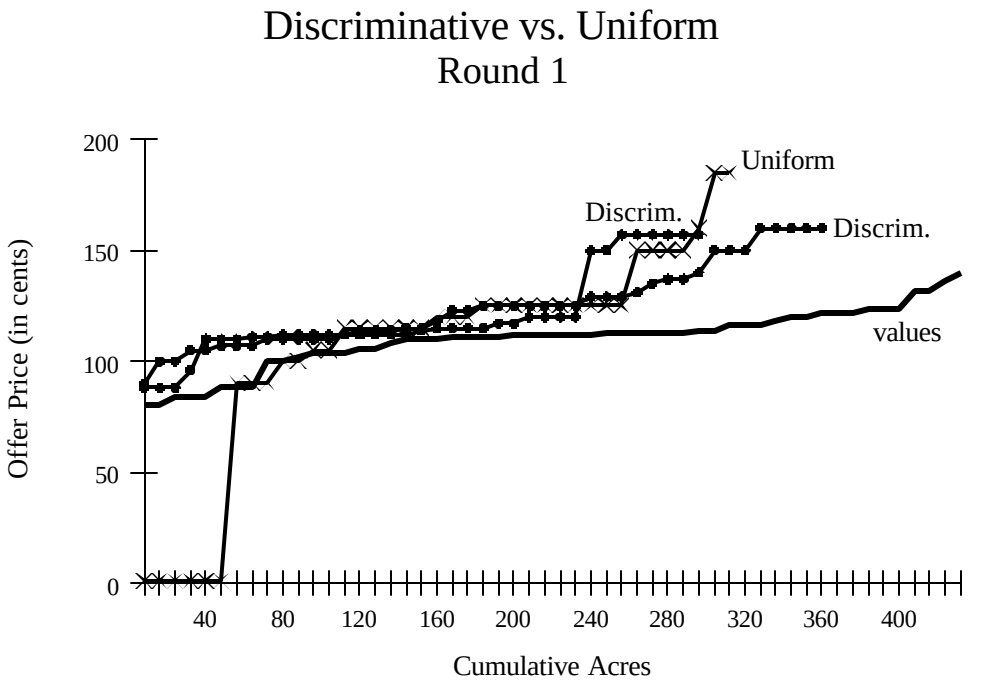

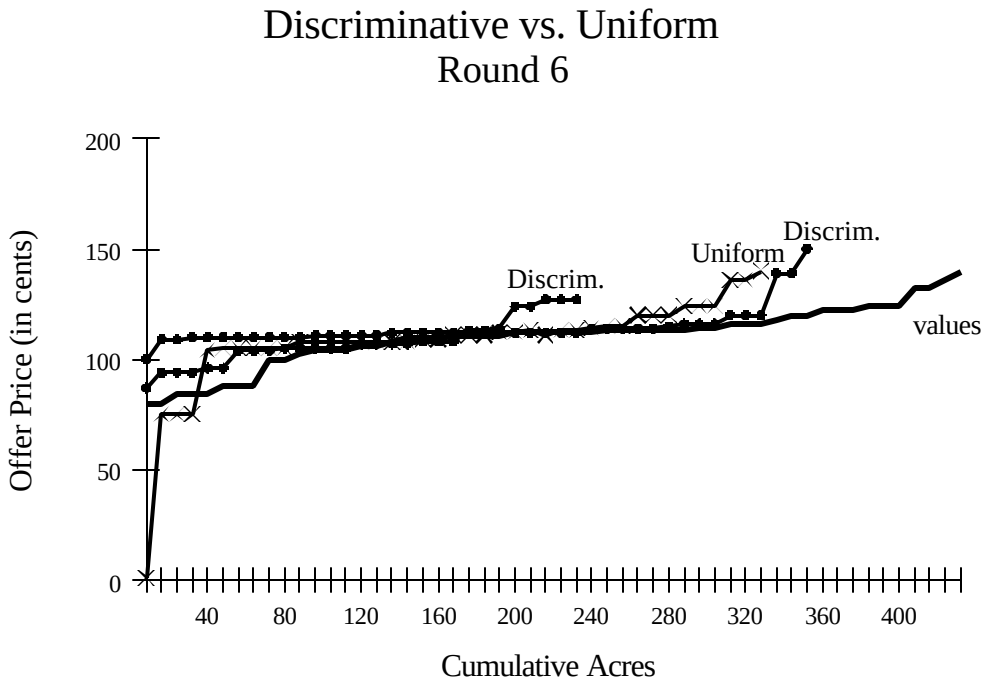

Figure 4. Average Price Paid in Uniform and Discriminative Price Auctions

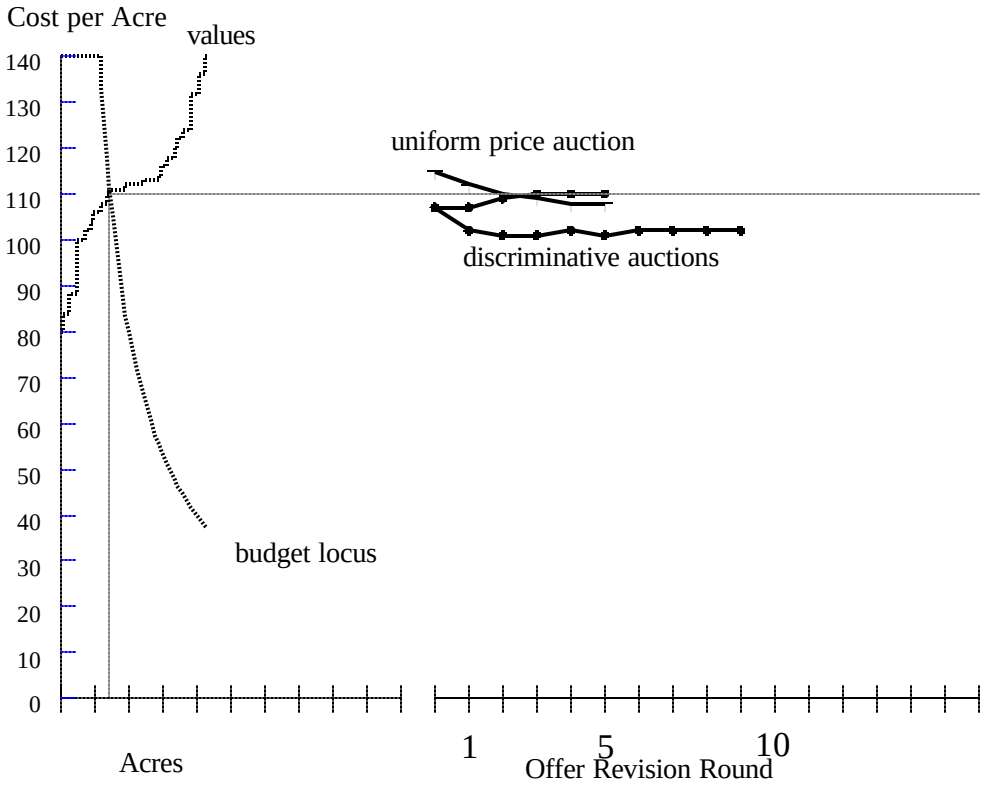

Figure 5. Average Price Paid in Discriminative Auctions: 20- and 42-subject auctions

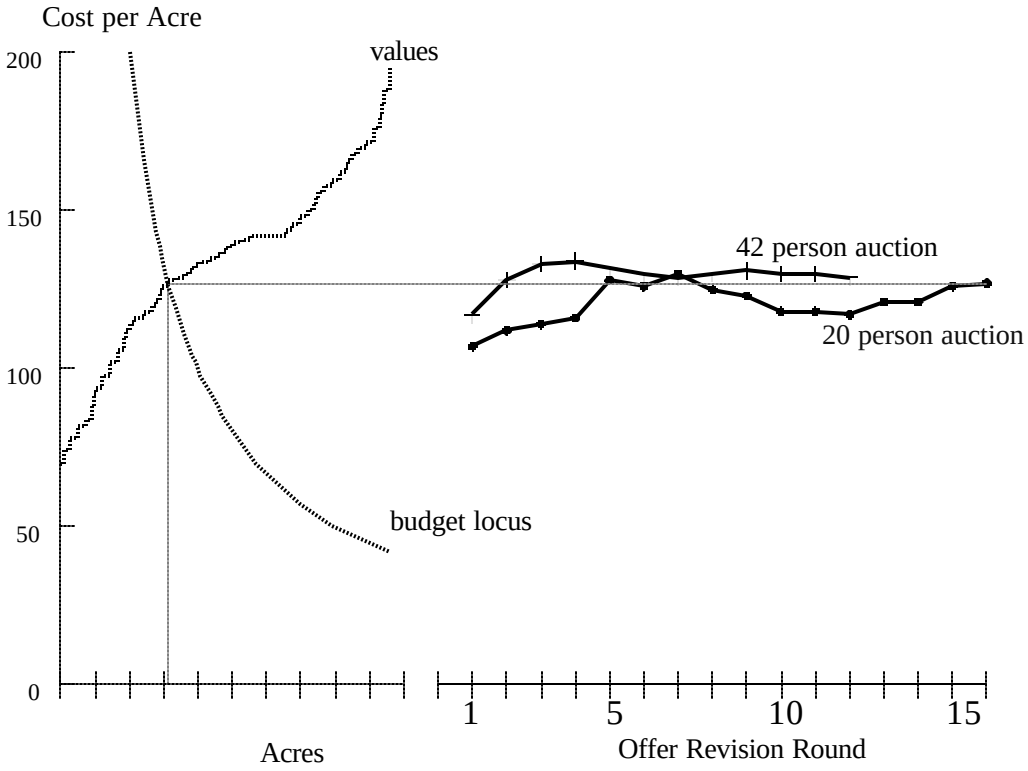

Figure 6. Time Series of Individual Offers in a Discriminative Auction
Key: Accepted Offers: Squares; Rejected Offers: Triangles

Offers: Subject 1

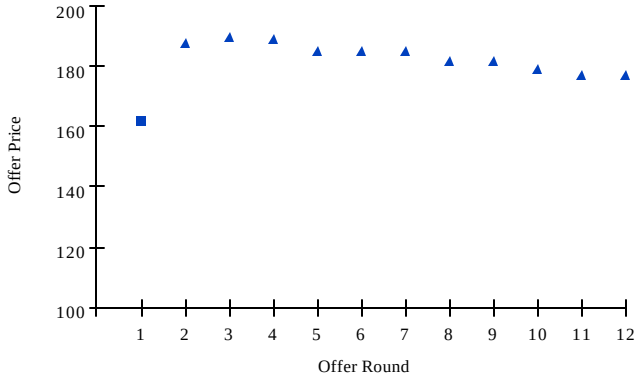

Offers: Subject 2

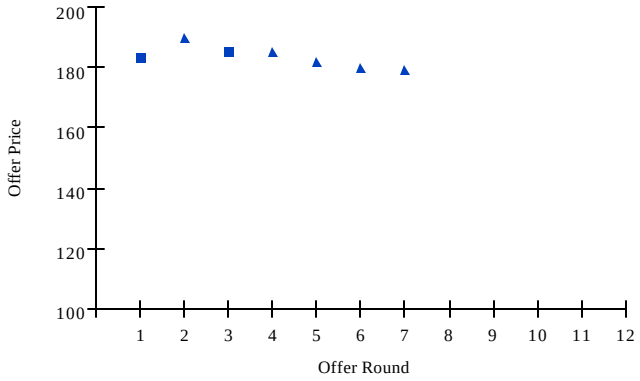

Offers: Subject 3

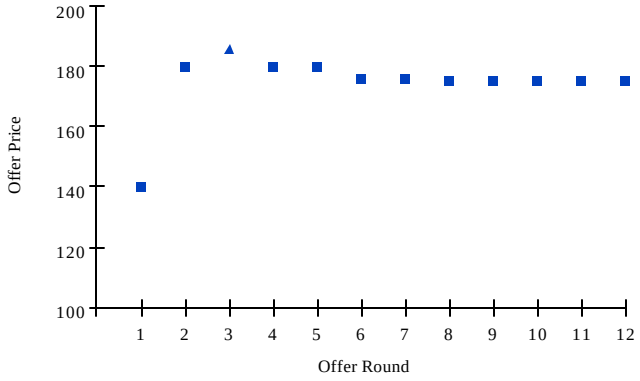

Figure 7. Average Price Paid in a Discriminative Auction: Highest Accepted Price is Not Announced

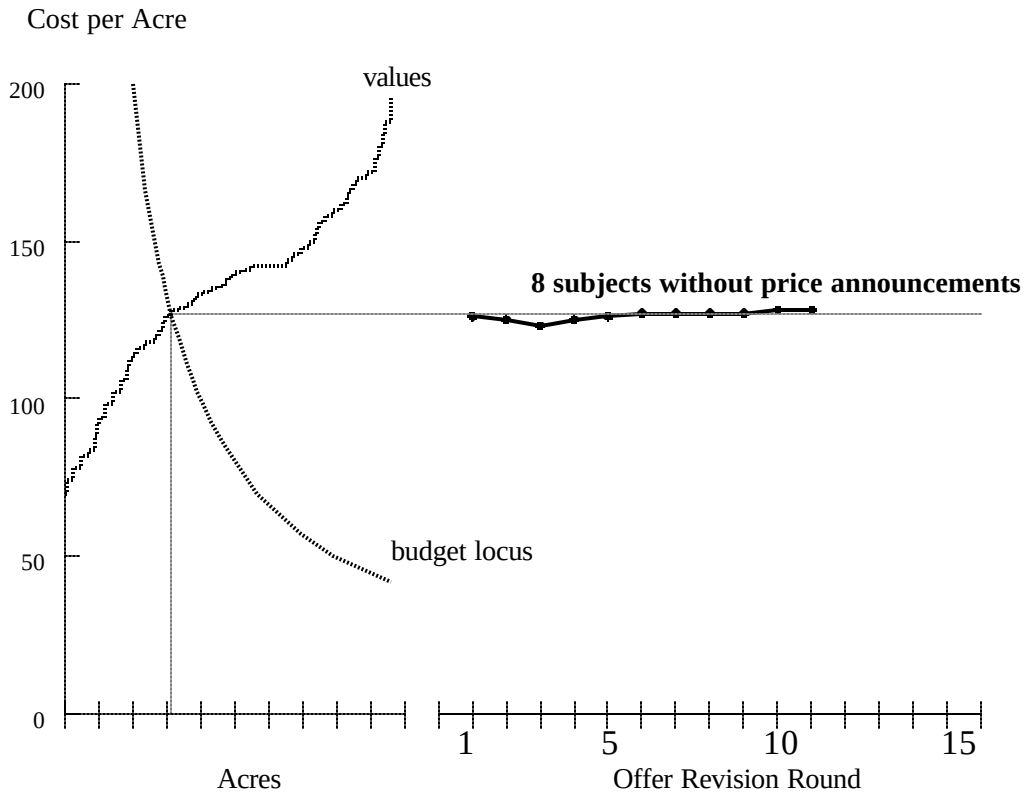

Figure 8. Comparing a One-Shot Discriminative Auction with a Discriminative Auction with Revisions

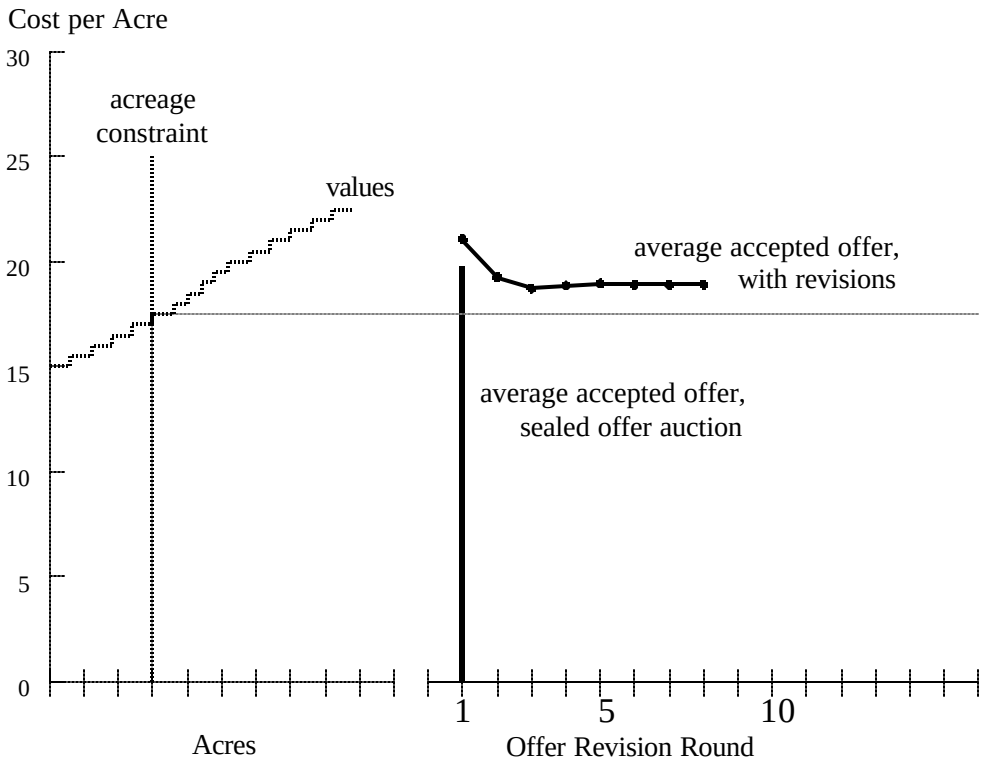

Figure 9. Offer Arrays from Field Test of Discriminative Auction Procedures

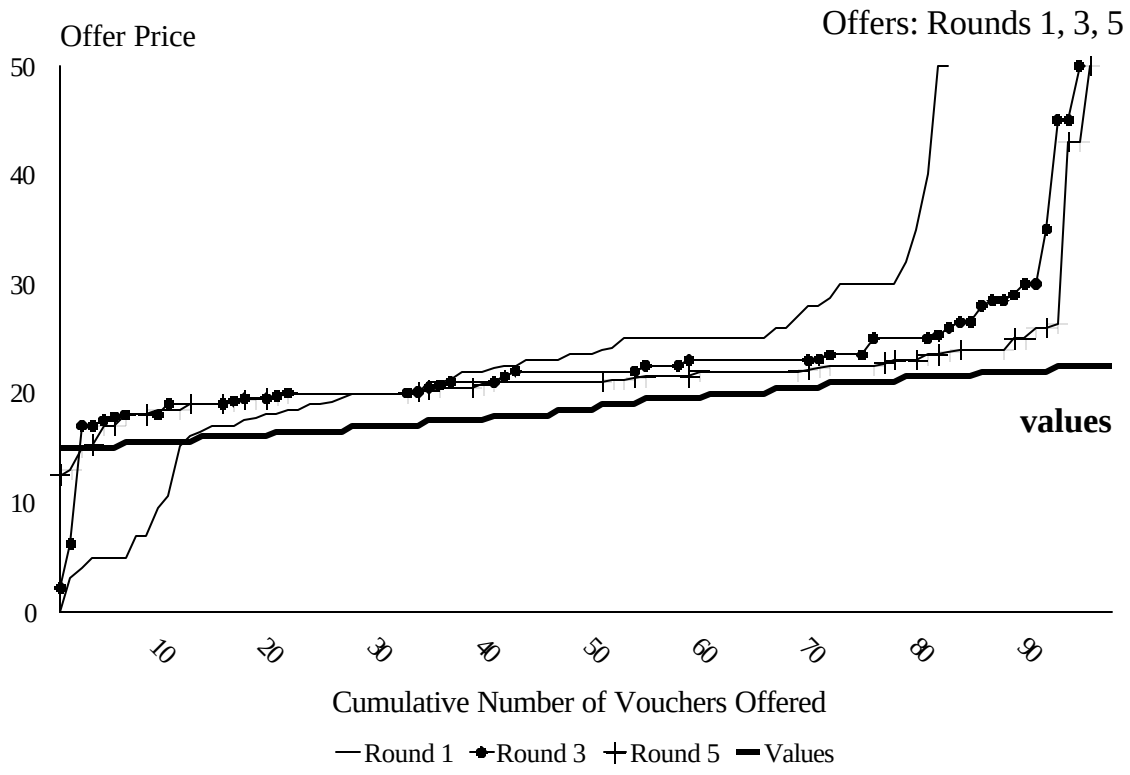

Figure 10. Offer Arrays from the EPD Irrigation Auction
Discriminative Auction Conducted on March 17, 2001

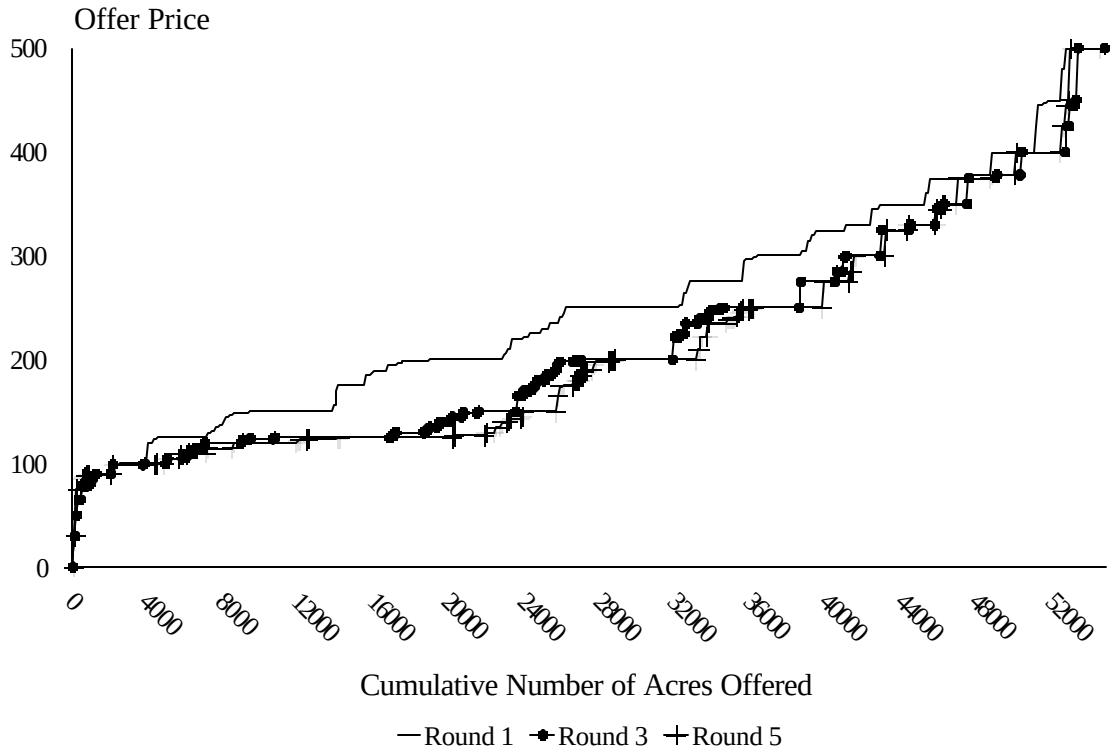

Figure 11. Offer Arrays from the EPD Irrigation Auction
Low, Medium, and High Offers