

DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES

CALIFORNIA INSTITUTE OF TECHNOLOGY

PASADENA, CALIFORNIA 91125

THE CMS AUCTION: EXPERIMENTAL STUDIES OF A MEDIAN-BID PROCUREMENT AUCTION WITH NON-BINDING BIDS

Brian Merlob

Charles R. Plott

Yuanzun Zhang



SOCIAL SCIENCE WORKING PAPER 1346

April 2011

The CMS Auction: Experimental studies of a median-bid procurement auction with non-binding bids

Brian Merlob

Charles R. Plott

Yuanzun Zhang

Abstract

We report on the experimental results of simple auctions with (i) a median-bid pricing rule and (ii) non-binding bids (winning bids can be withdrawn) - the two central pillars of the competitive bidding program designed by the Centers for Medicare and Medicaid Services (CMS). Comparisons between the performance of the CMS auction and the performance of the excluded-bid auction reveal the problematic nature of the CMS auction. The CMS auction fails to generate competitive prices for goods and fails to satisfy demand. In all proposed efficiency measures, we find the excluded-bid auction significantly outperforms the CMS auction.

JEL classification numbers:

Key words:

The CMS Auction: Experimental Studies of a Median-Bid Procurement Auction with Non-Binding Bids

Brian Merlob, Charles R. Plott and Yuanjun Zhang*
California Institute of Technology

On September 30, 2010, the *New York Times* published “[Fix Medicare’s Bizarre Auction Program](#),” which explained that “for the last ten years, the Centers for Medicare and Medicaid Services [CMS] has been testing an auction approach that is incredible in the inefficiency of its flawed design.” The controversy stems from two central features of the CMS auction that are not part of traditional auctions: (1) The price is set at the median of the winning bids, and (2) bids may be withdrawn after the price is announced. These two features make the CMS auction substantially different from traditional procurement auctions, and the theory of how the features interact has not been fully developed. In particular, the full structure and relationships of the Nash equilibria are unknown.

Critics of the proposed Medicare auction claim that the design encourages “low-ball” bids, in which suppliers submit the minimum allowable bid and then withdraw if the calculated market price is below their costs (Cramton and Katzman, 2010b). We conducted simple experiments to characterize the types of behavior and outcomes that could be expected from the proposed auction and find considerable evidence for these critics’ concerns. Auction architectures performing poorly in simple cases studied experimentally provide a realistic warning about problems that can surface in complex cases. Furthermore, if the behavior observed in the simple case is understandable in terms of theory, or even partial theory, then there is reason to take that theory seriously when applied to more complex cases.

The absence of a fully developed theory dictates an exploratory approach to experiments in which partial theories are tested and institutional comparisons are made. Through a comparison of institutions that are understood, we seek an understanding of the performance of the CMS auction. For comparison, we examine the theoretically understood excluded-bid auction. This auction has the same basic architecture as the CMS auction, but deviates by important features. In addition, we study institutions that lie “between” the CMS auction and the excluded-bid auction.

* The financial support of the Gordon and Betty Moore Foundation is gratefully acknowledged. Financial support for the sessions conducted jointly between Caltech and the University of Maryland was provided by an NSF grant to Peter Cramton and by the Gordon and Betty Moore Foundation. Kathryn Peters and Andre Pradhana contributed to the conduct and analysis of the experiments. The CMS auction was brought to our attention by Peter Cramton who has been central to our understanding of the auction and facilitated the experimental sessions conducted jointly between the Caltech Laboratory for Experimental Economics and Political Science (EEPS) and the University of Maryland Experimental Economics Laboratory. Our understanding of the Nash equilibria of the CMS auction benefited from the comments and writings of Brett Katzman. We thank Erkut Ozbay who organized the experimental laboratory at the University of Maryland and Kristian Lopez-Vargas who joined the experimental team for conducting the experiments during the joint Caltech and University of Maryland experimental sessions.

The basis of institutional comparisons is the performance of the CMS auction and other auctions in relation to four natural policy goals: (i) The auction should be successful in procuring the units demanded; (ii) The auction should be efficient from a social point of view in the sense that units are purchased from the lowest cost producers; (iii) The auction should not be wasteful from the government's point of view; and (iv) The auction should produce a competitive price that is capable of creating a healthy supply industry. Experiments explore alternative auctions that could be used in place of the CMS auction.

The report contains five sections plus appendices. Section 1 contains a description of the auction institutions. Section 2 describes the experimental design. Section 3 outlines central aspects of theory. In particular, key Nash equilibria are identified where they are known, and special features of bidding incentives are outlined in the cases for which the Nash equilibria are not known. Section 4 contains the results. The first set of results focuses on the excluded-bid auction and establishes the fact that it performs well with reasons that are understandable. The second set of results demonstrate that the CMS auction is not a good auction. The results reveal the existence and severity of its architectural flaws. The third set of results demonstrates that the CMS auction cannot be easily fixed. These results are acquired by experiments involving auctions with variations on the central pillars of the CMS auction. The conclusions drawn from these results suggest that simply removing a faulty procedure does not correct the problems with the proposed auction and may create new problems. The collections of results in Section 4 support the robustness of the conclusions and provide an intuitive theory about why these auctions have the properties discussed. Section 5 is a summary of the results.

1. INSTITUTIONS

Comparisons between the outcomes and behavior of the CMS auction with the excluded-bid auction provide an anchor for this study. While the CMS auction shares the same basic architecture as the theoretically robust excluded-bid auction, it deviates by important features. By perturbing these features in different experimentally controlled environments, we gauge the effects of particular design elements.

Each auction is a sealed-bid procurement auction, in which multiple units are sold at a uniform price to a single buyer with fixed demand. The institutions differ only by the pricing rule and by the bid withdrawal rule. In both institutions, bids are ordered from least to greatest. Then, a number of bids equal to the buyer's demand are accepted, and the corresponding suppliers are given the right to sell their unit at the market price. In the excluded-bid auction, the price is set at the value of the lowest bid that is not accepted. Furthermore, bids are binding commitments. In contrast, the price in the CMS auction is set at the median of accepted bids, and bids are non-binding commitments that may be withdrawn. Regardless of the number or pattern of bid

withdrawals, the procurement price is determined by the original set of winning bids in the auction.

In addition to pricing rule and bid withdrawal, two environmental variables were manipulated to determine robustness of the institutions. Increasing the number of bidders from twelve to sixteen people studied effects of scale. Under some conditions, bidders were shown no information about the costs of other bidders; in other conditions, bidders were given full information (costs and bids of other participants). Finally, a cost of participation was added in some experimental conditions in which bidders were required to pay a fee in order to place a bid.

In total, five features (each with two different states) were manipulated, offering thirty-two possible variations of the basic auction architecture. Of these possible conditions, only ten were studied based on their ability to reveal the most about the CMS auction in terms of changes in the efficiency and robustness in response to manipulation of key auction features. Figure 1 contains the relationship among the institutions and environments that were studied.

2: METHODS

2.1: Experimental Procedures

We conducted four experiments using identical experimental parameters. Each experiment involved different students recruited from laboratory subject databases at the California Institute of Technology and at the University of Maryland. All experiments were done on a computer terminal in an experimental economics computer lab at their respective universities under the supervision of the experimenters. Instructions were printed, distributed, and read to subjects. Following the introduction of each new auction, a short quiz was administered to check for comprehension. All costs were stated in experimental currency units (ECUs). Each experiment lasted 1.5 hours during which the number of participants was held constant.

A set of instructions are in Appendix B. The experimenters also used graphics to explain the structure of the auctions to participants. These graphics are in Appendix C.

2.2: Experimental Parameters

In all experiments, participants are sellers of a single unit with randomly assigned costs drawn from the uniform distribution [100, 1000]. The “government” is the only buyer and has a demand of seven units. The minimum allowable bid was 50 ECUs while the maximum allowable bid was 1050 ECUs.

The experiments at Caltech involved 12 participants each and were run on October 24, 2010 (experiment 20101024) and October 27, 2010 (experiment 20101027). Both experiments at the University of Maryland involved 16 participants each and were run in the morning (experiment 20101111-a) and afternoon of November 11, 2010 (experiment 20101111-b). Combined, the

experiments tested 10 different conditions listed in Table 1. For convenience, Table 2 lists each experiment with details on logistics, parameters, exchange rates, and the sequence of conditions used.

2.3: Experimental Design

Modifications of auctions in the experiments were either institutional or environmental. Due to limitations on time, only 10 out of the 32 possible conditions were tested. These 10 conditions are listed in Table 1 and were chosen because they reveal the most about the properties and robustness of the central pillars in the CMS auction.

The experiments are named by the date they were conducted. Experiments 20101024 and 20101027 contained six separate auctions, but looked at five conditions. The six auctions were (in order): the excluded-bid auction, the CMS auction with binding bids, the CMS auction, the CMS auction with full information, the CMS auction, and the excluded-bid auction with non-binding bids.

As a control, the excluded-bid auction was conducted first to establish a baseline for expected behavior in a theoretically understood auction. Subsequent auctions modified only one rule at a time to ease the transition between conditions for the participants. The CMS auction was conducted twice, with the full information CMS auction in between. This design exists to gauge the effects of learning for this auction mechanism. The experiments explore both central pillars (first the pricing rule, then the bid withdrawal rule). Furthermore, it also tests whether the results for an auction with well-known properties can hold when either of the pillars is imposed on it.

Experiments 20101111-a and 20101111-b were designed with the intent to determine whether results from the experiments at Caltech held given environmental changes. Both experiments at the University of Maryland planned for six auctions within the 1.5 hours. However, due to time constraints and difficulties in conducting long-distance experiments, only the first four were completed in Experiment 20101111-a and the first five in Experiment 20101111-b. The five auctions were (in order): the excluded-bid auction, the CMS auction, the CMS auction with full information, the excluded-bid auction with bidding fee, and the CMS auction with bidding fee (Experiment 20101111-b only).

3. THEORIES AND MODELS

A complete theory for the CMS auction and variations of the CMS auction does not exist. However, the theory that does exist is able to provide hints towards the effects of independent features of the auctions. We start with the excluded-bid auction, which has a well-defined Nash equilibrium before turning to the CMS auction, in which we identify one family of Nash equilibria in our experimental environment. Because the CMS auction has a family of equilibria,

it lacks uniqueness, which leads to some ambiguity about which equilibrium, if any, is a foundation for the observed behavior. Nevertheless the family of equilibria provides insights to the behaviors observed. The following discussion focuses on variations of the CMS and excluded-bid institutions and the currently existing theories of behavior.

3.1 *The Excluded-Bid Auction with Binding Bids*

The excluded-bid auction has market price determined by the first excluded bid. All bids below the market price are accepted and receive the market price for the unit. Bids are binding in the sense that accepted bids must be sold at the market price. Because the excluded-bid auction is well understood, it was chosen to serve as a reference point. It produces competitive revenue, encourages suppliers to reveal their true costs, and assigns goods efficiently. The auction satisfies *voluntary participation* because at the Nash equilibrium bidders are guaranteed a non-negative payoff.

When suppliers have only one unit to sell, the auction has a single Nash equilibrium. The auction institution is a truthful mechanism—every bidder maximizes his payoff by bidding his true cost, regardless of the other bids submitted. In theory, this auction should achieve consistently high efficiencies regardless of the experience of the bidders, the extent of information available, or the scale of application.

3.2 *The Median-Bid Auction with Non-Binding Bids (CMS)*

The CMS auction has two key features: the price is determined by the *median winning bid*¹ and *bids are non-binding*. This auction is not truthful—at the Nash equilibrium suppliers misrepresent their costs. However, this auction does satisfy voluntary participation because suppliers may withdraw their bids if the calculated market price is below their cost.

The CMS auction has multiple Nash equilibria, including the extreme case of every supplier submitting the minimum allowable bid (Cramton and Katzman, 2010b). For one family of Nash equilibria in the CMS auction, a number of bidders greater than the median of the number of demanded units ($n > m$) bid the minimum allowable, thereby setting the price at the minimum allowable bid. No single bidder can improve his payoff by submitting a different bid, so it constitutes a Nash equilibrium.

Interestingly, despite this family of equilibria, all bidders have a weakly dominating strategy to submit a bid above their cost and withdraw if the market price is below their cost. To prove this, suppose that a supplier submits a bid below or equal to his cost ($b \leq c$). Assuming he wins (otherwise, his payoff is zero), either (1) the price will be below or equal to his cost ($p \leq c$) and

¹ For k demanded units, the median is defined as $(k+1)/2$ if k is odd and $k/2+1$ if k is even.

his payoff will be at most zero ($p - c \leq 0$) or (2) the price will be above his cost ($b \leq c < p$) and his payoff will be positive. However, whenever the price is above his cost, he would have achieved equal profits by bidding his cost (because it would not affect the median) and greater profits by bidding above it. Thus, a supplier maximizes his profits by bidding above his cost.

When costs and bids are common knowledge (public information) in the CMS auction, a different class of Nash equilibria emerges. Let $n(B)$ be the set of bidders with costs at most B . Choose any bid B for which $|n(B)| \leq k$, for k demanded units. Let all bidders in the set $n(B)$ bid B , and all other bidders bid at least B (and withdraw if they win). In these equilibria, the price is set at B and no single bidder can change the price.

It is of interest to note that another equilibrium exists in which the bidders in $n(B)$ bid B and withdraw their bids if the resulting price is below their cost while all other bidders bid the minimum allowable and withdraw.

An observation about non-binding bids is useful, especially in the context of the CMS auction. If a bid can be withdrawn and if the bid does not determine the price, such as in median-bid price auctions, then the bid only works to determine whether the bidder is among the winners or not. If the price is below the bidder's cost, the bidder can withdraw his bid. A low-ball bid that ensures the bidder will be among the winners is an option that gives the bidder the right to sell at the market price. While exercise of the withdrawal option is not part of any equilibrium strategy that we know, bidders do choose it (as illustrated in Section 4). One speculative theory suggests that the bidders have malicious intent to hurt the other bidders by forcing the price down. To test that theory we add a cost to bid in two of the experiments. If malicious intentions exist, then this behavior will continue at some price in the form of the low-ball bids. If on the other hand, the low-ball bids reflect a naïve strategy, the addition of the bidding fee will remove it from behavior.

3.3 The Excluded-Bid Auction with Non-Binding Bids

One variation of the CMS architecture changes the pricing rule from the median price rule to the excluded bid rule. This produces an auction with the excluded bid pricing rule and non-binding bids. Similar to CMS auction, this variation satisfies voluntary participation because participants have the option to withdraw bids. Similar to the excluded-bid auction, this auction has an equilibrium bid function $bid(cost) = cost$. However, due to the withdrawal option, the family of equilibria identified for the CMS auction also extends to this auction: if a number of bidders greater than the number of demanded units ($n > k$) bid the minimum allowable, then the price is set at the minimum and no single bidder can improve his payoff by submitting a different bid. Furthermore, this variation is not a truthful mechanism because a bidder's optimal bid depends on all other submitted bids.

3.4 The Median-Bid Auction with Binding Bids

A second variation of the CMS architecture keeps the median-bid pricing rule and removes the bid withdrawal possibilities, thereby changing the rule to binding bids. Similar to the CMS auction, this auction is not a truthful mechanism; participants have an incentive to bid above their cost. For this auction, Cramton and Katzman (2010b) show that there are no strictly increasing bid functions. Instead, they suggest that suppliers with high costs will attempt to avoid negative payoffs by bidding the maximum allowable bid. Suppliers cannot be guaranteed non-negative payoffs regardless of their bids, so this auction does not satisfy voluntary participation.

3.5 Bidding Fees

The addition of a bidding fee encourages bidders who expect to be excluded from the winners to opt out of the bidding process. In the excluded-bid auction with binding bids, if fewer than eight bids were tendered the price was set at the institutional maximum known to all.

Under the Nash equilibrium of the excluded bid rule, bidders opt out only if they had costs over 575 ECUs. While they only have a 10.5% chance of winning (at 575 ECUs), when bidders do win, they receive a large payoff (425 ECUs). Consequently, we predict that 16.1% of trial periods will have less than perfect procurement. It remains a truthful auction, but the dominant strategy shifts to

$$bid(cost) = \begin{cases} cost + fee, & \text{if } \sum_i^{k-2} \binom{n-1}{i} cost^i (1-cost)^{n-1-i} > fee \\ opt\ out, & \text{otherwise} \end{cases}$$

for #suppliers $n = 16$, and units demanded $k = 7$, $cost \approx 575$

In the CMS auction, the introduction of a bidding fee creates an incentive for those who expect to be excluded to opt out of the bidding process. However, the Nash equilibrium is unknown to us since suppliers who do not expect to be excluded may still choose to opt out if they believe that the median cost will be below their cost plus the bidding fee.

Section 4: RESULTS

The research is summarized by the following 10 results. Results 1 – 3 collectively argue that procurement auctions satisfying CMS’s four policy goals exist and can be compared to the CMS auction mechanism. In particular, the excluded-bid auction performs well in the sense that the performance is understandable from a theoretical point of view, efficient, and satisfies the natural policy objectives. Results 4 – 6 reveal, however, that the CMS auction does not perform as

satisfactorily. Results 7 and 8 report on the individual features that make the CMS auction unusual. Results 9 and 10 explain the effects of adding a bidding fee.

RESULT 1: In the excluded-bid auction, suppliers tend to reveal their costs.

Support. Figures 4 and 10 display the bidding behavior for the excluded-bid auction. Theory suggests that it is a dominant strategy for suppliers to bid their costs. Indeed, three out of every four bids (78.1%) approximately² follow this strategy in Experiments 20101024 and 20101027. As we increase the number of suppliers in Experiments 20101111-a and 20101111-b, the adherence to this strategy drops to 54.0% and 68.3% respectively. A third of suppliers consistently submit bids equal to their costs while the remaining suppliers, on average, tend to overbid when they have low costs. This observation is captured by a simple linear regression of the model $bid = \alpha + \beta * cost$, which reveals a relatively small intercept $\hat{\alpha} = 123$ ($se = 12.5$) and a slope $\hat{\beta} = .856$ ($se = .020$). These parameters also suggest that suppliers are more likely to bid truthfully as their costs approach the expected competitive price (which is equal to the cost of the excluded bidder).

RESULT 2: Prices generated in the excluded-bid auction approximate the competitive price.

Support. In the excluded-bid auction, 22 out of the 52 total periods generate prices within 1% of the competitive price; 36 out of 52 are within 5%. As depicted in Figures 2 and 3, the prices that differ greatly from the competitive price are located both above and below the competitive price and result largely from untruthful bidding by one or two bidders (who either received negative payoffs or lost the right to sell). Consequently, the average market price generated is 102.2% ($se = 13.1$) of the competitive price. Similar to the model for bidding behavior, a simple linear regression for the model $Market Price = \alpha + \beta * (Competitive Price)$ returns an intercept $\hat{\alpha} = 111$ ($se = 35.5$) and a slope $\hat{\beta} = .824$ ($se = .057$) that is significantly greater than zero, with a p-value $< 10^{-18}$. If we constrain α to 0, then $\hat{\beta} = 1.00$ ($se = .014$), closely matching the ideal $\beta = 1$.

RESULT 3: Outcomes in the excluded-bid auction are consistent with policy objectives: the auction has perfect procurement with near-perfect efficiency levels.

Support. As is true for all binding bids auctions with sufficient participation, the excluded-bid auction has perfect procurement in every period. From the government perspective, perfect procurement at prices that approximate the competitive price (as reported in Result 2) lead to an efficiency level of 100.3% ($se = 14.8\%$). Efficiency from the supplier perspective is more

² We define “approximately” as ± 30 ECUs (10€) to take into account the tendency of participants to round their cost to the nearest 5, 10 or 25. To give some perspective, in Experiments 20101024 and 20101027, 40.0% follow the dominant strategy exactly, 58.7% are within ± 2 ECUs, and 70.8% are within ± 15 ECUs.

variable due to untruthful bidding, but still high on average at 93.7% ($se = 42.3\%$). Social efficiency – how well the auction allocates units to suppliers with the lowest costs – is near-perfect at 96.1% ($se = 4.7\%$).

Results 4 – 6 suggest that the CMS auction performs poorly in absolute terms and relative to the excluded-bid auction. The CMS auction fails to provide a sufficient disincentive for perverse bidding behavior that leads to low efficiency and poor performance.

RESULT 4: In the CMS auction, suppliers with high costs submit “low-ball” bids. The prevalence of this activity increases with experience, information, and scale.

Figures 5 and 11 display the bidding behavior for the CMS auction with private information; Figures 8 and 9 display the bidding behavior for those with public information. In each figure, there is a collection of bids at the minimum allowable level for suppliers with primarily high costs. Across all CMS auctions, one out of every five bids is a low-ball bid. The prevalence of this behavior more than doubles with experience (from 12% to 26%), and is exacerbated by information (to 24%), and scale (to 18%). In Experiments 20101024 and 20101027, many of the remaining bids are scattered above the supplier’s cost, as suggested by the incomplete theory presented in Section 3. There is less of a discernable trend in the CMS auctions of Experiments 20101111-a and 20101111-b. Simple linear regressions of the model $bid = \alpha + \beta * cost$ are not appropriate because they describe the data poorly and thus are not reported. However, if we ignore low-ball bids in Experiments 20101024 and 20101027, we get an intercept $\hat{\alpha} = 207$ ($se = 11.8$) and $\hat{\beta} = .805$ ($se = .021$) for private information, and $\hat{\alpha} = 271$ ($se = 25.7$) and $\hat{\beta} = .712$ ($se = .047$) for public information.

RESULT 5: Prices generated in the CMS auction do not approximate the competitive price and are significantly and consistently lower than the competitive price.

Support. Figures 2 and 3 show that the CMS auction generates prices that are consistently lower than the competitive price – only 5 out of the 138 total periods generate prices within 15% of the competitive price. The average market price generated differs greatly with experience, information, and scale. For 12 suppliers and private information in CMS Iteration I, the average price is 69.5% ($se = 19.5\%$) of the competitive price, which drops with the introduction of public information to 39.2% ($se = 29.8\%$) and drops slightly further with experience to 34.9% ($se = 25.2\%$) in CMS Iteration II. Scaling up to 16 suppliers, the average price begins low at 33.4% ($se = 18.6\%$) for private information and rises to 46.1% ($se = 23.1$) for public information. This information is summarized in Table 3. For each CMS auction, the econometric model $Market Price = \alpha + \beta * (Competitive Price)$ returns a slope $\hat{\beta}$ that is not statistically significantly greater than 0, suggesting that underlying competitive price has no bearing on the observed market price.

RESULT 6: Outcomes in the CMS auction are not consistent with policy objectives: the auction fails to obtain adequate procurement and has poor efficiency levels relative to the excluded-bid auction

Support. In the CMS auction, one out of five periods procured no units at all and only 3 of the 138 total periods had perfect procurements. On average, the CMS auction with 12 bidders and private information procured 4.2 units ($se = 1.7$) out of 7 units demanded in Iteration I, which declined to 2.2 units ($se = 2.2$) for Iteration II. As the auction increased its scale to 16 bidders, the average dropped to 1.7 units ($se = 1.8$). When valuations were public information, an average of 2.2 units ($se = 2.4$) and 2.4 units ($se = 1.8$) were procured for 12 and 16 bidders respectively. A consequence of low prices and poor procurements, efficiency from the supplier perspective suffered dramatically with 29.3% at best and 1.1% at worst. Although the government benefited from low prices, inadequate procurements led to efficiency levels (73.9%, 52.0%, 37.2%, 47.5%) that were poor relative to the near-perfect efficiency reported in Result 2. Social efficiency in the CMS auction is also much lower (53.8%, 35.3%, 25.3%, 37.1%) than the in the excluded-bid auction because not all units are procured. See Table 3 for more details.

Results 7 and 8, suggest that the CMS auction cannot be fixed by marginal changes. The problems exist at the foundations of the auction design. The policy of non-binding bids can independently make an otherwise well-functioning auction perform poorly. Similarly, the median-bid pricing rule has the independent ability to disrupt well-functioning auctions.

RESULT 7: In the excluded-bid auction with non-binding bids, suppliers submit “low-ball” bids, causing lower prices, procurement, and efficiencies than those observed in the excluded-bid auction.

Support. Similar to low-ball bids reported in Result 4 in the CMS auction, Figure 6 shows that suppliers in the excluded-bid auction with non-binding bids submit bids at the minimum allowable level. The remaining suppliers tend to bid above their cost, perhaps accounting for the low-ball bidders. Figure 2 depicts that the resulting market prices are often lower than the competitive price, but greater than the prices observed in the CMS auction and reported in Result 5. On average, the price is 83.9% ($se = 17.9\%$) of the competitive price. Out of the 24 total periods, 9 have perfect procurement, 5 procure six, 7 procure five, and the remaining 3 procure four units. Thus, this auction averages 5.8 units. Efficiency from the perspective of government is 112.0% ($se = 30.5\%$) because the government is procuring all or nearly all demanded units at prices lower than the competitive price. Consequently, the low prices and low procurements cause efficiency from the supplier perspective to be 68.5% ($se = 32.9\%$), lower than in the excluded-bid auction. Social efficiency remains high at 87.9% ($se = 11.5\%$).

RESULT 8: *In the median bid auction with binding bids, suppliers consistently bid above their cost. This auction procures the demanded number of units at prices below the competitive price at high levels of efficiency for the government, but at low levels of efficiency for suppliers.*

Support. As shown in Figure 7, suppliers in median-bid auction with non-binding bids submit bids above their costs. In contrast to the auctions with non-binding bids reported on in Results 4 and 7, no supplier with a high cost submits a low-ball bid. Instead, suppliers are bidding at the maximum allowable level, as predicted by the incomplete theory. A simple linear regression of the model $bid = \alpha + \beta * cost$ reveals an intercept $\hat{\alpha} = 170$ ($se = 19.7$) and a slope $\hat{\beta} = 1.00$ ($se = .032$). These parameters suggest that, on average, suppliers submit bids that are 170 greater than their cost. As shown in Figure 2, these bids translate into prices that approximate the competitive price better than the CMS auctions but are still lower than the competitive price on average at 88.7% ($se = 13.7\%$). However, binding bids and sufficient (in this case, mandatory) participation implies perfect procurement, even if suppliers are forced to sell at a loss (at prices below their bids). This leads to above-perfect levels of efficiency from the government perspective at 124.6% ($se = 30.2\%$), near-perfect social efficiency at 97.4% ($se = 3.1\%$), but poor efficiency for suppliers at 68.6% ($se = 34.6\%$).

Results 9 and 10 address the speculative theory that bidders are acting maliciously by imposing a bidding fee and observing whether low-ball bidding behavior persists.

RESULT 9: *The imposition of bidding fees eliminates low-ball bidding behavior. Instead, suppliers with high costs opt out of the bidding process.*

Support. Figure 12 shows the bidding behavior in the excluded-bid auction with a bidding fee. Theory suggests that suppliers bid truthfully until a certain cost and then opt out. A simple linear regression of the model $bid = \alpha + \beta * (cost + fee)$ returns an intercept $\hat{\alpha} = 142$ ($se = 20.8$) and a slope $\hat{\beta} = .711$ ($se = .044$), revealing little evidence for truthful bidding. Suppliers do appear to start opting out at the right cost though - starting at costs of about 600 (in comparison to 575 for theory) suppliers are more likely to opt out than bid. Figure 13 shows the bidding behavior in the CMS auction with a bidding fee. Theory suggests that suppliers bid above their cost, and indeed, a simple linear regression of the model $bid = \alpha + \beta * (cost + fee)$ returns an intercept $\hat{\alpha} = 48.5$ ($se = 14.3$) and a slope $\hat{\beta} = .912$ ($se = .048$), indicating a progression towards revealing costs. In this auction, however, suppliers start to opt out at lower costs – only 7 bids are made after costs of 400.

RESULT 10: *The excluded-bid auction with a bidding fee approximates the competitive price better than the median bid auction with a bidding fee, resulting in better procurements and higher levels of efficiency from the perspective of the suppliers or society, but not from the perspective of the government.*

For these two auctions, prices are on the opposite extremes: the excluded-bid auction has average prices at 120% ($se = 32.4\%$) of the competitive price as opposed to the CMS auction with average prices at 60.3% ($se = 15.8\%$) of the competitive price. The high prices in the excluded-bid auction are due primarily to the rule that the price is set at the maximum allowable level if less suppliers participate than the number of demanded units plus one (i.e., if there is no excluded bidder). Ignoring these instances returns an average price just slightly above the competitive price at 104.7% ($se = 19.2\%$). As a consequence of the price, average procurement in the excluded-bid auction with a bidding fee remains high (relative to the other excluded-bid auctions) at 6.8 units ($se = .4 \text{ units}$) while average procurement in the CMS auction with a bidding fee remains low (relative to the other CMS auctions) at 3.57 units ($se = 1.6 \text{ units}$). High procurement and high prices lead to high supplier efficiency at 162.5% ($se = 77.0\%$) and near-perfect social efficiency at 95.5% ($se = 5.8\%$) in the excluded-bid auction while low procurement and low prices lead to low supplier efficiency at 34.4% ($se = 23.0\%$) and low social efficiency at 59.9% ($se = 19.3\%$) in the CMS auction. From the perspective of government, however, the occasional price at the maximum allowable level in the excluded-bid auction brings efficiency down to 66.5% ($se = 33.7\%$), but ignoring these instances brings it back up to 84.9% ($se = 15.4\%$). The government doesn't fare much better in the CMS auction, with an efficiency of 74.6% ($se = 24.9\%$).

5. SUMMARY AND CONCLUSION

The central issues explored in this paper are motivated by the architecture of the CMS auction and by the policy goals the auction was designed to address. An assessment of the performance is based on an institutional comparison of the CMS auction with its variations and on applications of theory. Evidence suggests that the withdrawal option contributes to the presence of low-ball bidding behavior and that the median-bid pricing rule reliably generates prices lower than the competitive price, deterring suppliers from submitting bids when participation is not mandatory (Result 10). Hence, the CMS auction rests on an architecture that constrains its efficiency below the levels at which existing, well-known auctions are capable of performing.

In the classical uniform price auction, the price is determined by the first excluded bid and bids are binding commitments. Comparing this excluded-bid auction with the CMS auction under controlled, laboratory conditions reveals the superior performance of the excluded-bid auction. The excluded-bid auction performs as predicted by theory: the prevalent bidder strategy is to reveal costs (Result 1), and as a consequence, prices emerge near the competitive equilibrium price (Result 2) and the social efficiency is nearly perfect. (Result 3) Furthermore, the auction behavior remains consistent with theory regardless of scale or experience.

In contrast to the excluded-bid auction, the CMS auction performs poorly as a procurement auction. A pervasive bidding strategy is to bid at the minimum allowable level (Result 4), driving

down prices (Result 5) and generating a low efficiency on each dimension (Result 6). Furthermore, the poor performance of the CMS auction is exacerbated by experience, information, and scale.

Analysis of select features from the CMS auction architecture help isolate independent contributions to the poor performance and suggest that “slight” adjustments to the CMS auction will not improve performance to the standards set by the excluded-bid auction. In particular, replacing the median-bid pricing rule with the first-excluded-bid pricing rule (while maintaining the right to withdraw) allows for the independent study of the impact of the withdrawal option (Result 7). The resulting increase in number of equilibria introduced by the withdrawal option removes the opportunity to produce any “crisp” test of theory, but reveals the option’s effect on bidding behavior. In general, suppliers were observed to either place bids near their cost (except those with low costs, who safely bid above their cost without risking their right to sell) or at zero, two of the equilibria identified (Result 7). Removal of the bid withdrawal feature while retaining the pricing rule dramatically (but predictably) increased prices and procurement, but the prices observed were reliably lower than the prices in the excluded-bid auction and contributed to low supplier efficiency. With the bid withdrawal feature removed, bidders must trade at the price set by the median bid, which may be substantially below their cost. Thus, attempting to avoid such a result, there was evidence that bidders with costs above the expected market price attempted to exclude themselves by submitting a high bid (Result 8).

The imposition of a bidding fee eliminated low-ball bidding (Result 9), suggesting that bidders are not malicious, but rather that auction architecture does not provide sufficient disincentives for underbidding. Furthermore, for suppliers with medium to high costs, the CMS auction deters participation because the prices generated do not approximate the competitive price (Result 10).

The use of experiments to develop an understanding of what might take place in more complex and uncontrolled field environments relies substantially on theory. The chain of logic is made difficult in the case of the CMS auction because the theory is incomplete. Nevertheless, partial theories contribute to our understanding of particular observed behaviors and institutional comparisons allow us to assess the performance of novel auction architectures. The challenge, for those who are unconvinced by the evidence that the CMS auction will perform poorly, is to explain why the principles seen operating so clearly in the laboratory environment will not operate when the auction is deployed in the field.

REFERENCES

- Ayres, I. and P. Cramton (2010). "Fix Medicare's Bizarre Auction Program". *New York Times*.
- Cason, T. N. and C.R. Plott (1996). "EPA's New Emissions Trading Mechanism: A Laboratory Evaluation". *Journal of Environmental Economics and Management*, 133-160.
- Cramton, P. and B. Katzman (2010a). "Reducing Healthcare Costs Requires Good Market Design". *The Economists' Voice*.
- Cramton, P. and B. Katzman (2010b). *Designed to Fail: The Medicare Auction for Durable Medical Equipment*. Working Paper, University of Maryland.
- Kagel, J. H. and D. Levin (2001). "Behavior in Multi-Unit Demand Auctions: Experiments with Uniform Price and Dynamic Vickrey Auctions". *Econometrica*, 69(2): 413-454.
- Krishna, V. (2002). *Auction Theory*. San Diego: Academic Press.
- Miller, G. J. and C.R. Plott (1985). "Revenue Generating Properties of Sealed-Bid Auctions: An Experimental Analysis of One-Price and Discriminative Processes". In V. Smith (Ed.), *Research in Experimental Economics* 3:159-182). Greenwich, CT: JAI press Inc.

Appendix A: Figures and Tables

Figure 1:

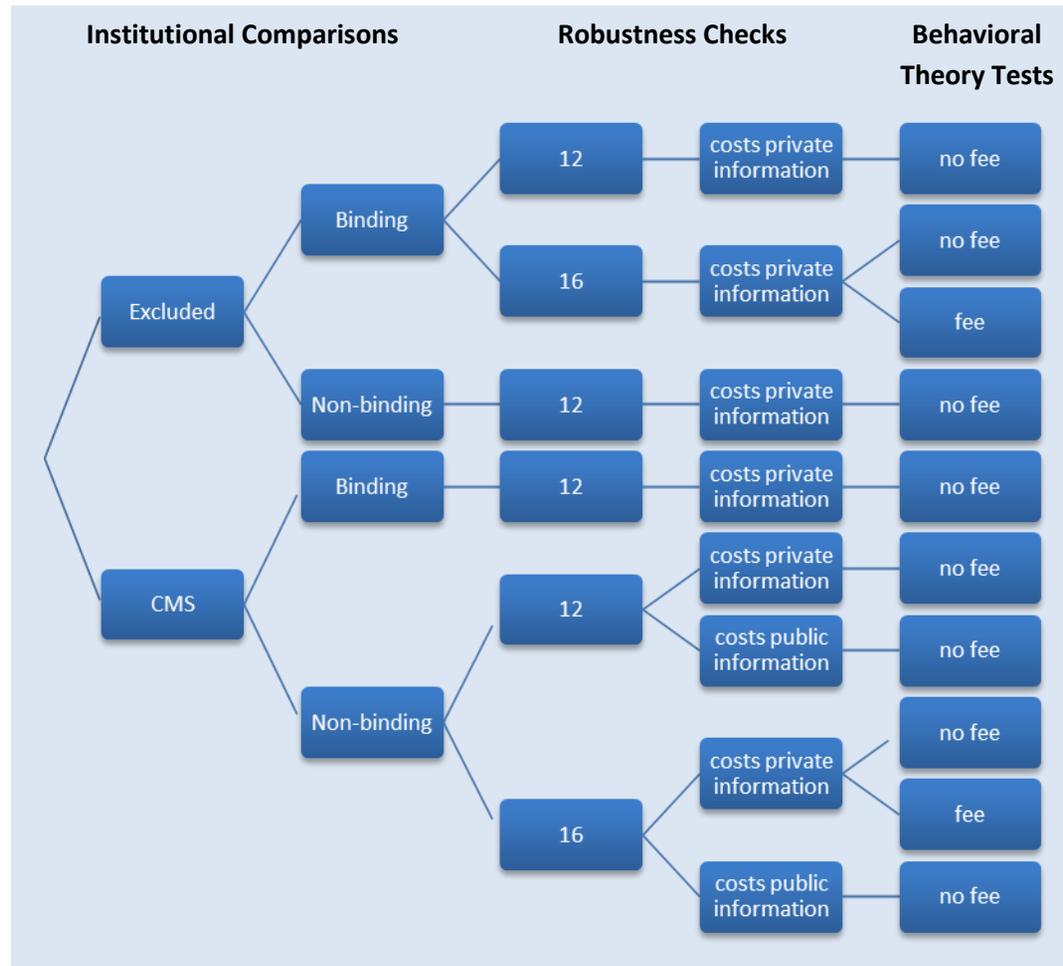


Figure 2: Experiments 1 & 2: Pricing Success

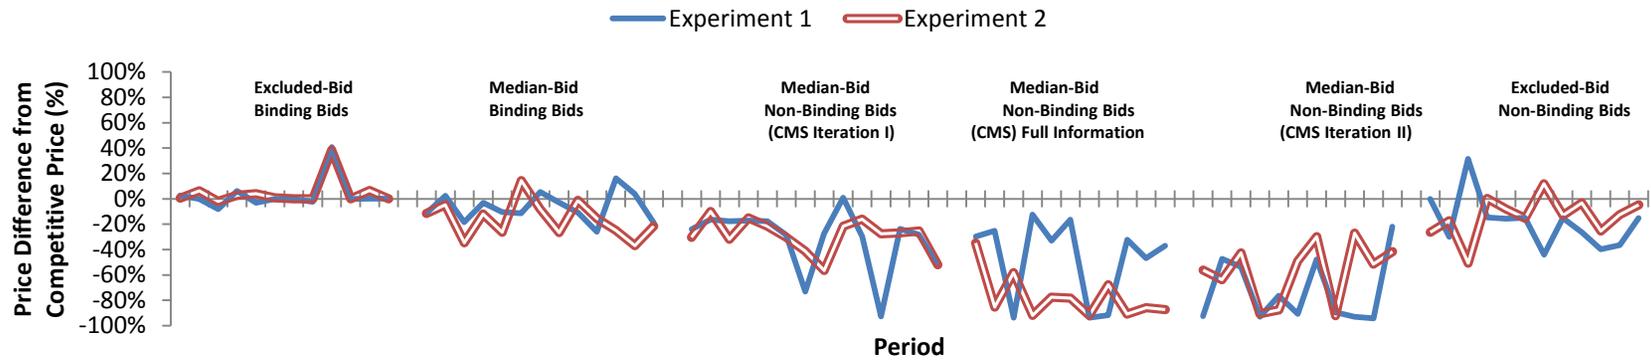


Figure 3: Experiments 3 & 4: Pricing Success

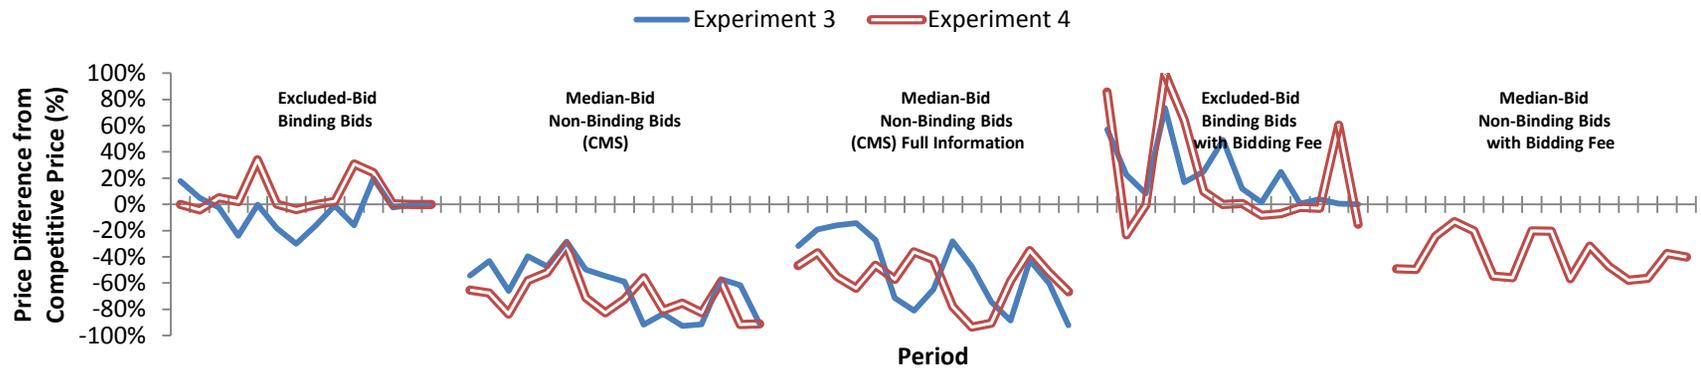


Figure 4: Excluded-Bid Auction with Binding Bids

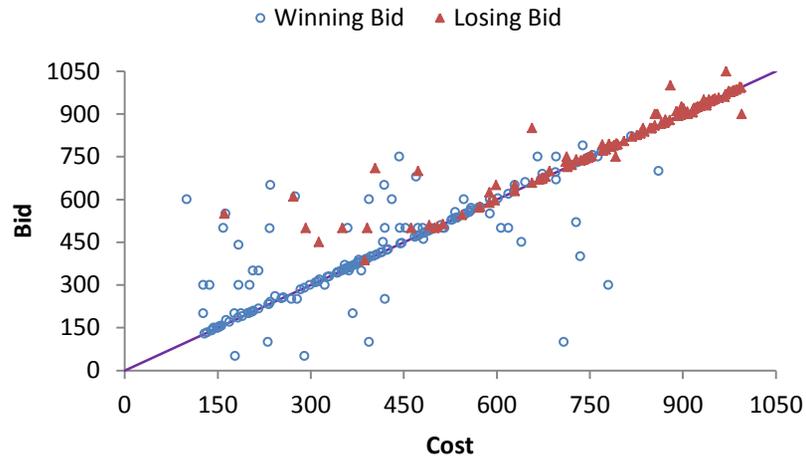


Figure 5: Median-Bid Auction with Non-Binding Bids (CMS)

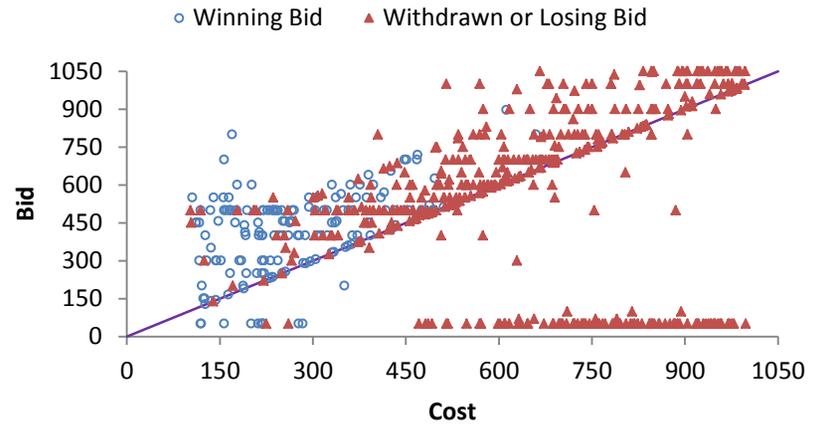


Figure 6: Excluded-Bid Auction with Non-Binding Bids

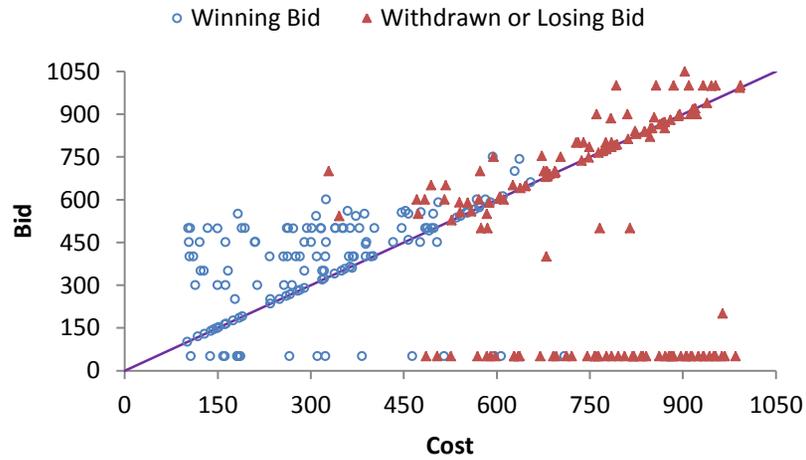


Figure 7: Median-Bid Auction with Binding Bids

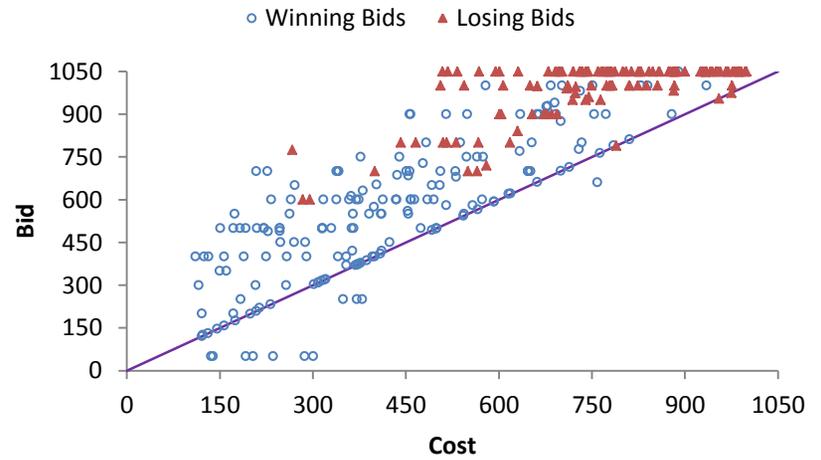


Figure 8: Median-Bid Auction with Non-Binding Bids (CMS): Full Information

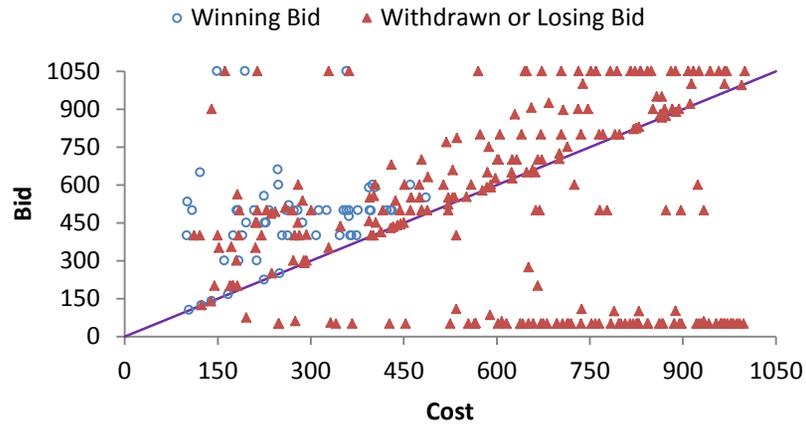


Figure 9: Median-Bid Auction with Non-Binding Bids (CMS): Full Information and 25% More Bidders

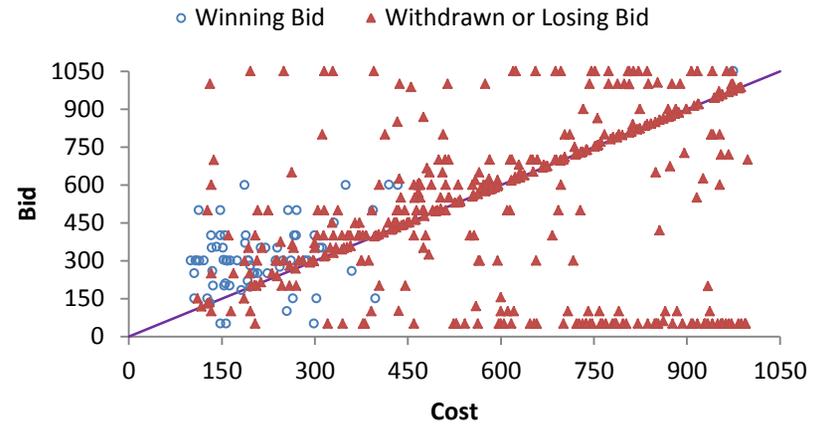


Figure 10: Excluded-Bid Auction with Binding Bids: 25% More Bidders

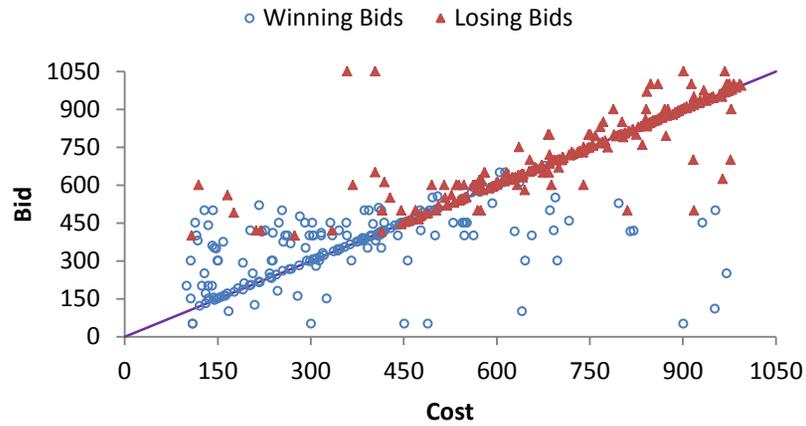
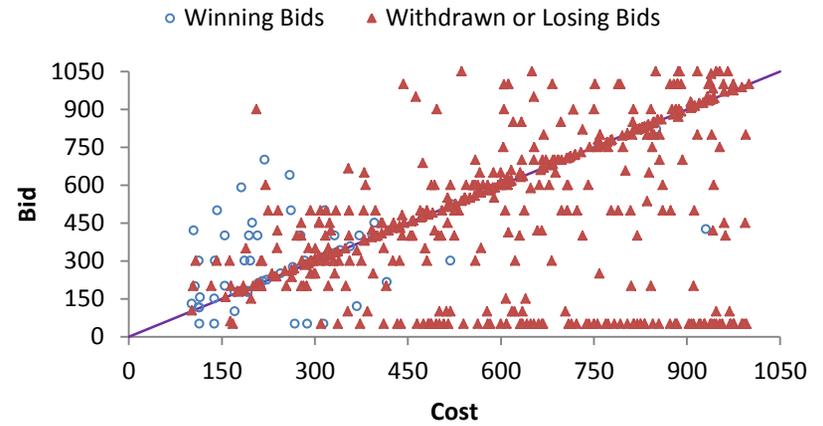
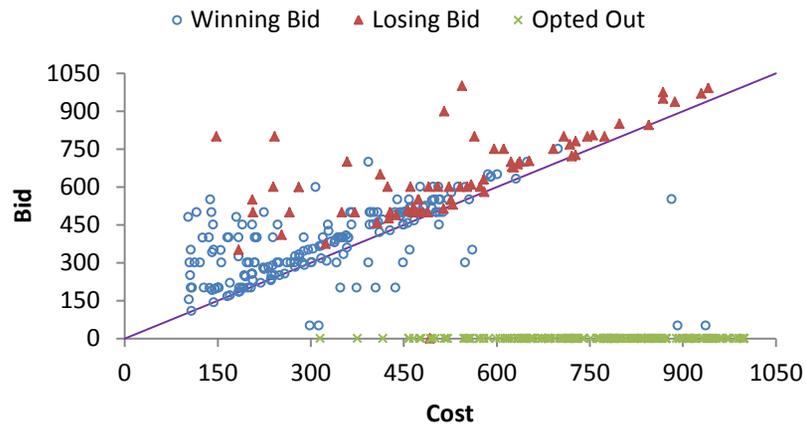


Figure 11: Median-Bid Auction with Non-Binding Bids (CMS): 25% More Bidders



**Figure 12: Excluded-Bid Auction with Binding Bids:
Bidding Fee and 25% More Bidders**



**Figure 13: Median-Bid Auction with Non-Binding Bids
(CMS): Bidding Fee and 25% More Bidders**

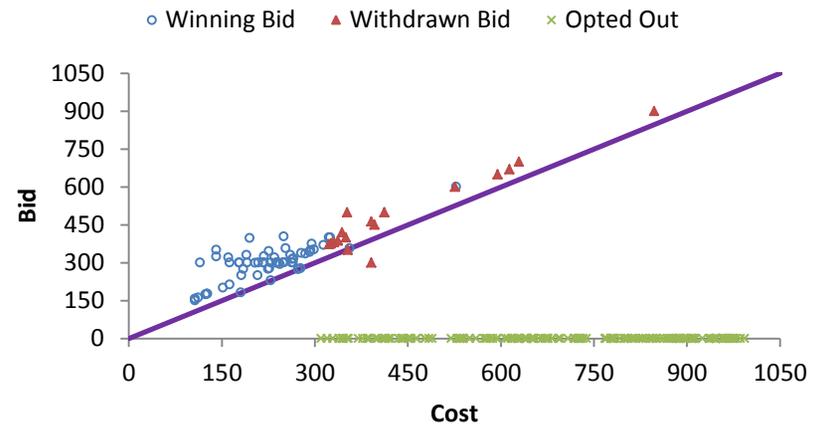


Table 1

Condition	Pricing Rule	Bid Withdrawal Rule	Number of Participants	Nature of Information	Bidding Fee
1	Excluded-Bid	Binding	12	Private	No
2	Excluded-Bid	Binding	16	Private	No
3	Excluded-Bid	Binding	16	Private	Yes
4	Excluded-Bid	Non-Binding	12	Private	No
5	CMS	Binding	12	Private	No
6	CMS	Non-Binding	12	Private	No
7	CMS	Non-Binding	12	Public	No
8	CMS	Non-Binding	16	Private	No
9	CMS	Non-Binding	16	Private	Yes
10	CMS	Non-Binding	16	Public	No

Table 2

Experiment #	1	2	3	4
Date	10-24-2010	10-27-2010	11-11-2010	11-11-2010
Location	California Institute of Technology		University of Maryland	
Exchange Rate (ECUs/\$)	300		125	
Conditions Run (in order)	1, 5, 6, 7, 6, 4		2, 8, 10, 3	2, 8, 10, 3, 9

Treatment	Procurement (units)		Pricing Success		Social Efficiency		Government Surplus		Supplier Surplus	
	Mean	SE	Mean	SE	Mean	SE	Mean	SE	Mean	SE
Excluded-Bid (12)	7	0	103.9%	11.4%	97.7%	4.2%	96.9%	10.1%	101.6%	31.3%
Excluded-Bid (16)	7	0	100.8%	14.6%	94.7%	4.9%	103.2%	14.8%	86.9%	49.5%
Excluded-Bid (16), Fee	6.66	.45	104.7%	19.2%	95.5%	5.8%	66.5%	37.8%	162.5%	77.0%
Excluded-Bid (12), Non-B	5.83	1.09	83.9%	17.9%	87.9%	11.5%	112.0%	30.5%	68.5%	32.9%
CMS (12), Binding	7	0	88.7%	13.7%	97.4%	3.1%	124.6%	30.2%	68.7%	34.6%
CMS (12), Public Info.	2.19	2.40	39.3%	29.8%	35.3%	36.8%	52.0%	54.2%	17.8%	23.3%
CMS Iteration I (12)	4.21	1.77	69.5%	19.5%	56.2%	24.8%	70.8%	31.4%	42.3%	23.3%
CMS Iteration II (12)	2.23	2.22	34.9%	25.2%	36.6%	35.2%	61.1%	69.4%	12.8%	15.7%
CMS (16), Public Info.	2.37	1.77	46.1%	23.1%	37.1%	26.5%	47.5%	33.7%	13.4%	21.4%
CMS (16), Fee	3.75	1.57	60.3%	15.8%	59.9%	19.3%	74.6%	24.9%	34.4%	23.0%
CMS (16)	1.72	1.84	33.4%	18.6%	25.3%	26.8%	37.2%	40.9%	0.7%	8.3%

Appendix B: Instructions

Excluded Bid:

In this auction, you are going to participate in **an excluded bid auction** in which you will be a SELLER of a commodity. In this auction with 12 participants, the participants submitting the lowest 7 bids will sell the item at the price of the 8th lowest bid.

You are selling a single item. Your cost of production for the single item is drawn from a uniform distribution [100, 1000]. Your task is to submit bids for the opportunity to sell your item. In each period you will be randomly assigned a different cost of production. You are not to reveal this information to anyone. It is your own private information. The minimum bid you may submit is 50; the maximum bid you may submit is 1050.

$$Profit = \begin{cases} \text{if you submit one of the 7 lowest bids:} & (\text{market price}) - (\text{cost of production}) \\ \text{otherwise:} & 0 \end{cases}$$

$$\text{market price} = 8^{\text{th}} \text{ lowest bid}$$

Questions

Suppose for example, the bids are: 100, 120, 130, 140, 150, 160, 170, 180, 190, 200, 210, and 220.

- What is the highest accepted bid? _____
- What is the price that all bidders will receive? _____
- If Participant A's bid is 160 and his cost of production is 185. What is his profit? _____
- If you bid *below* your value, can you lose money? Yes No
- If you bid your value, can you lose money? Yes No
- If you bid *above* your value, can you lose money? Yes No

Median Price Auction with Cancellation:

Participants with the 7 lowest bids will be offered the right to sell the item at the valued price. When you are offered the right to sell your item, you can also opt to 'cancel your bid' and not sell the item at the market price. A bid cancellation *will not* affect the valued price of the item. You will all choose to accept or cancel simultaneously, but if a cancellation occurs, then the person with the next highest bid who has accepted the market price will sell the item.

$$Profit = \begin{cases} \text{if you one of the 7 lowest bidders that **accept**:} & (\text{market price}) - (\text{cost of production}) \\ \text{otherwise:} & 0 \end{cases}$$

$$\text{market price} = \text{median of 7 lowest bids}$$

Median Price Auction with Cancellation [with public information]:

However, in this auction, you will be able to see the costs of all participants. Once all participants have bid, you will also be able to see their bids and whether they are one of the lowest 7 bidders. The minimum bid you may submit is 50; the maximum bid you may submit is 1050.

Entry Fee:

For the remaining auctions, after learning your cost, you will first decide whether to bid in the particular auction. **To bid in the auction, you must pay an entry fee of 50 ECUs.** Thus, your profit in Experimental Currency Units (ECUs) will be determined by:

$$\text{Your profit} = \begin{cases} \text{if you do not bid:} & 0 \\ \text{if you sell the item:} & (\text{market price}) - (\text{cost of production}) - 50 \\ \text{if you don't sell the item:} & - 50 \end{cases}$$

Entry Fee and Excluded Bid:

In this auction, you are going to participate in an **excluded bid auction** in which you will be a SELLER of a commodity. **To bid in this auction, you must pay an entry fee of 50 ECUs.** Bidders submitting the lowest 7 bids will sell the item at the price of the 8th lowest bid. If there are fewer than 7 bidders, the item will be sold at 1050 ECUs.

$$\text{market price} = \begin{cases} \text{if } > 7 \text{ bidders:} & 8^{\text{th}} \text{ lowest bid} \\ \text{if } \leq 7 \text{ bidders:} & 1050 \end{cases}$$

CMS with Entry Fee:

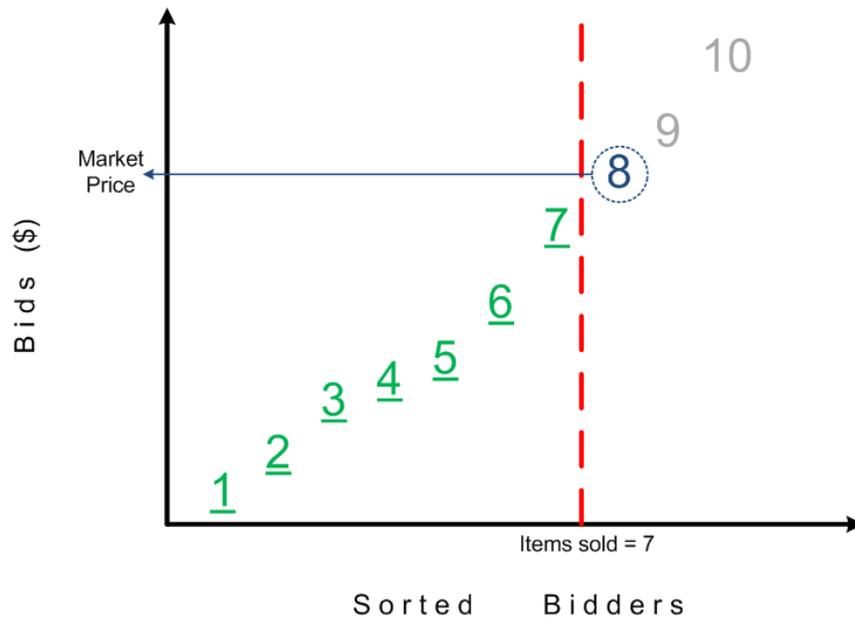
In this auction, you are going to participate in a **median price auction with cancellation** in which you will be a SELLER of a commodity. **To bid in this auction, you must pay an entry fee of 50 ECUs.** After all bidders submit bids, the median of the lowest 7 bids will be calculated and **set irrevocably** as the market price. If there are fewer than 7 bidders, then the market price is **set irrevocably** at the median among all submitted bids (if there are an even number of bidders, the price is the higher of the two middle bids).

$$\text{market price} = \begin{cases} \text{if } \geq 7 \text{ bidders:} & \text{median of 7 lowest bids} \\ \text{if } < 7 \text{ bidders:} & \text{median of bids} \end{cases}$$

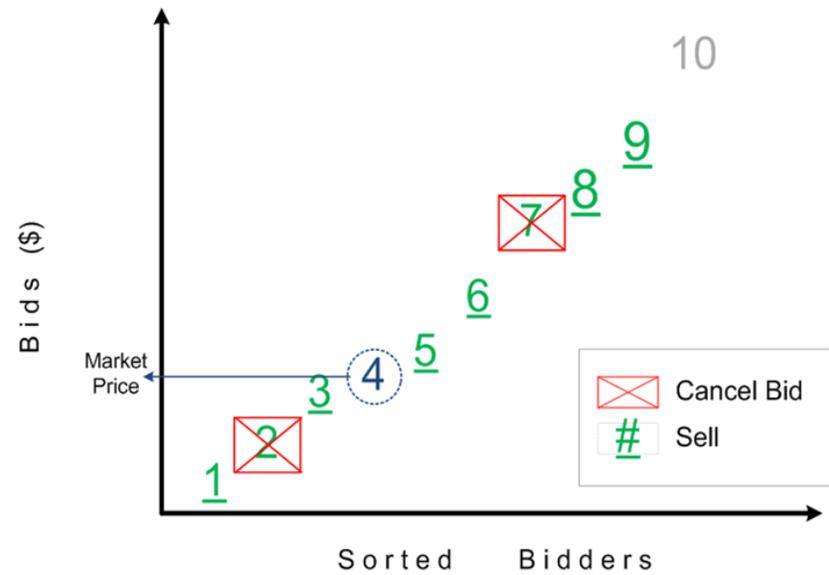
Appendix C: Graphics

The following were used as visual aids when explaining the instructions in the experiments. Images at Caltech were hand-drawn on the white board, while those at the University of Maryland were produced by Peter Cramton and are shown below:

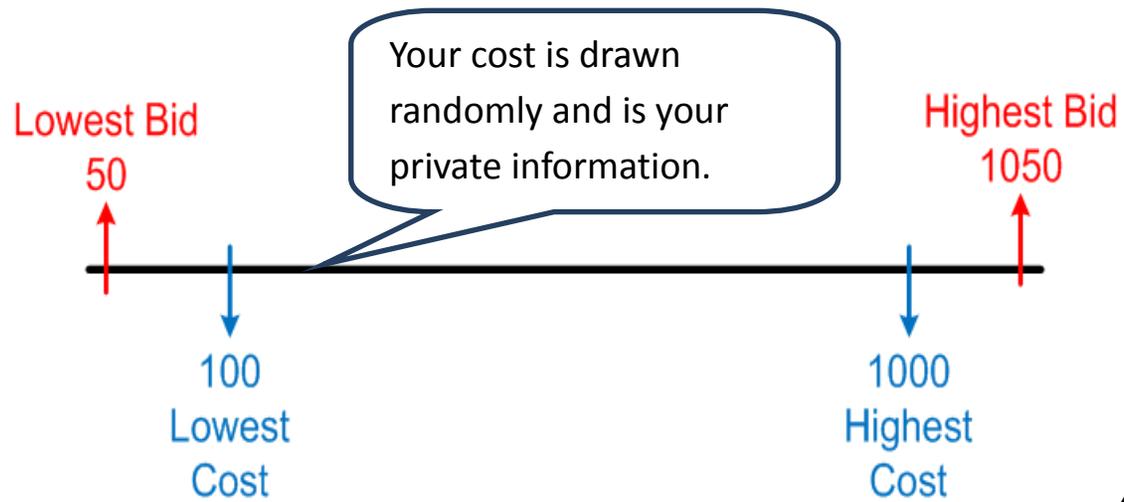
FIRST EXCLUDED BID



MEDIAN PRICE WITH CANCELLATION



Possible Costs and Admissible Bids



4