

of HAWAI'I°

MĀNOA

# University of Hawai`i at Mānoa Department of Economics Working Paper Series

Saunders Hall 542, 2424 Maile Way, Honolulu, HI 96822 Phone: (808) 956 -8496 www.economics.hawaii.edu

Working Paper No. 14-9

Using Quotas to Enhance Competition in Asymmetric Auctions: A Comparison of Theoretical and Experimental Outcomes

By

Daniel Hellerstein, Nathaniel Higgins, Michael J. Roberts

March 2014

# Using Quotas to Enhance Competition in Asymmetric Auctions: A Comparison of Theoretical and Experimental Outcomes

Daniel Hellerstein<sup>a</sup>, Nathaniel Higgins<sup>a</sup>, Michael J. Roberts<sup>b</sup>

 <sup>a</sup>USDA Economic Research Service, 1800 M Street NW, Washington DC 20036 Email: danielh@ers.usda.gov, nhiggins@ers.usda.gov
 <sup>b</sup>Department of Economics & Sea Grant, Saunders 542, 2424 Maile Way University of Hawai'i at Manoa, Honolulu HI 96822 Email: michael.roberts@hawaii.edu

#### Abstract

We study multiple-unit asymmetric procurement auctions wherein sellers from two classes draw costs from different distributions. When sellers are asymmetric, a cost-minimizing buyer discriminates among classes of sellers to enhance competition [1]. Establishing a quota—a limit on the number of offers that can be accepted from any one class—discriminates simply and effectively. The quota increases demand scarcity from the perspective of lowcost sellers, which causes them to lower their offers. To solve for approximate equilibrium strategies of asymmetric auctions with and without a quota, we develop a new method that is similar but distinctly different from the constrained strategic equilibrium (CSE) approach [2]. The new method finds the constrained strategies that minimize the expected gain from a randomly chosen seller unilaterally deviating from the constrained strategy. We find quota can enhance competition and lower total procurement cost. We subject the same auctions to laboratory testing and find savings from quota in excess of that predicted by the approximate equilibrium strategies. This study is first to combine theory and experimental evidence of auctions with quotas, though similar mechanisms are widely used in practice. Because the mechanism is widely used to promote social goals and can also lead to better outcomes for the buyer, our findings have both positive and normative implications. One potentially interesting application of quota auctions would be for large-scale procurement of ecosystem services like carbon sequestration.

*Keywords:* asymmetric auction, optimal auction, experimental auction, multi-unit auction *JEL:* D44, C63, C91

#### 1. Introduction

When sellers are asymmetric, the optimal procurement auction is one that discriminates between sellers [1, 3]. While the conditions that characterize an optimal auction have been known for awhile, implementation remains an open issue. How can a buyer, knowing that sellers are observably different in their ability to produce a good, practically structure an auction that exploits this asymmetry to save procurement costs? Direct implementation of the optimal auction requires knowledge of the distribution of seller costs and the ability to discriminate perfectly between sellers. Since neither of these two conditions is likely to be met in practice, we investigate the returns to a simple mechanism that a buyer can easily implement. The mechanism imposes a quota, a limit on the number of winning offers that can come from any single class of observably similar sellers, to increase competition within that class. That is, the auctioneer specifies *ex ante* that he will accept no more than x offers from any defined class of sellers.

We find motivation to study simple price discriminating mechanisms in many practical applications. Private firms engage in contract procurement using auctions. Every level of government procures goods from suppliers that are observably heterogeneous in some way. Popular examples include defense-related procurement, procurement of infrastructure contracts, and procurement of fleet vehicles. The federal government also procures environmental services from heterogeneous private landowners using an auction procedure. <sup>1</sup> A similarly intriguing application could be for procurement of carbon offsets from highly heterogeneous land owners. A mechanism that encourages competition among similar types could also exist in markets with less structure than a formal auction. Firms, for example, hire from heterogeneous labor pools and might limit hiring from any particular pool so as to pay some workers less than their marginal productivity [5].

The quotas we examine are similar to what some in the auction literature call "set asides." When multiple units of a good are being auctioned by the government (whether items to be sold or contracts to be purchased), set-asides reserve some number to be won by qualified bidders. Qualified bidders are bidders selected based on observable characteristics, often race or business size, meant in most cases to promote social goals, such as encouraging participation by a minority class of bidders [6]. Milgrom [7] presents a simple example of how set-asides can increase competition in an auction and so enhance the auction outcome from the buyer's perspective.<sup>2</sup> The fact that set-asides are used both in the sale of public

<sup>&</sup>lt;sup>1</sup>The largest single example of a land conservation auction is the Conservation Reserve Program, implemented by the U.S. Department of Agriculture. Large-scale conservation auctions like this one would be especially suitable for the quota mechanism because many offers are selected from a large pool of observably heterogeneous sellers. Classes of sellers could be defined by location, size, or other characteristics. See Kirwan, Lubowski, and Roberts [4] for more detail about the program and estimates of gains to participating landowners.

<sup>&</sup>lt;sup>2</sup>Milgrom's example is simple in that distributions of bidder reservation values in a purchase auction do not overlap across classes of buyers. Thus, without set asides or bid preferences, bidders from the high-value class always win. Our examples are more general, as we allow cost distributions to overlap.

goods such as spectrum [5], and in government procurement, provides a positive motivation for our study. That is, in addition to or despite possible distributional goals, governments or firms may use quotas or set-asides simply to reduce procurement costs.

Quotas reduce procurement costs when sellers of several dissimilar classes compete to sell multiple goods. Sellers compete against rivals both similar and dissimilar to themselves.<sup>3</sup> Sellers from these dissimilar classes, having observable characteristics that distinguish them, will offer toward a common margin. This margin is set by a mix of within-class and between-class competition. When one class of sellers has lower opportunity costs than another, a quota enhances within-class competition. The intuition is straightforward: by limiting the number of winning offers, demand from that group declines. The artificial scarcity makes offers more competitive. The tradeoff is that between-class competition is sacrificed: when a low-cost seller is eliminated, high-cost sellers face less competition.

McAfee and McMillan [8] provide an example in a context of international trade, which we modify slightly. There are six firms, two foreign and low-cost and four domestic and highcost, competing for two government contracts. Unrestricted competition is characterized by weak competition within the class of low-cost foreign firms. The marginal foreign firm competes with domestic firms to fulfill the second contract, while the stronger foreign firm extracts substantial rents. McAfee and McMillan investigate how price preferences influence the procurement cost of an auction. If a quota were imposed that mandated a maximum of one foreign and one domestic firm to fulfill the governments need, the low-cost foreign firms would be forced to compete directly with each other. Rent that would have been extracted by low-cost foreign firms is reduced while rent accruing to domestic firms increases. The net effect of a quota depends on the net balance of offsetting influences: low-cost foreign firms face tougher competition, while high-cost domestic firms inflate their offers in the absence of direct foreign competition. In this polar example, quota has effectively created two separate auctions, one in which only foreign firms compete, and one in which only domestic firms compete.

The total effect of a quota on procurement cost is the sum of enhanced competition within classes and reduced competition between classes. A quota is thus most beneficial to the buyer when within-class competition is low among low-cost sellers. This happens if demand for the marginal unit typically comes from a high-cost class. In single-price auction without the restriction of a quota, what we refer to as an open auction throughout the paper,

<sup>&</sup>lt;sup>3</sup>Similarly, in an auction to sell rather than procure multiple items, a seller would benefit from quota when buyers of many different classes compete to purchase.

sellers from the low-cost class will offer toward the same margin as sellers from the high-cost class and thereby extract substantial rents. Setting a quota effectively reduces the surplus captured by low-cost sellers.

While the idea of discriminating among sellers was laid out in the seminal paper on optimal auctions [1], the method of discrimination considered here is new in a key way. Earlier papers have examined what have been called "bid preferences," which amount to a discounting of some offers relative to others for the purpose of determining winners. For example, auctions for the procurement of transportation contracts in California select a qualified small" contractor if its offer is within 5% of the lowest offer by an unqualified seller. [9, 10]. Such a preference could be formulated as a discount to the qualified offer for only the purpose of ranking [11]. That is, the buyer ranks discounted offers, equal to  $(1 - \text{preference}) \times \text{offer}$ , from lowest-to-highest, with a larger preference for small bidders and then selects the lowest discounted offer but pays the winner their full undiscounted offer. Such offer preference programs are common in government procurement and have been applied in high stakes auctions such as the first auctions for spectrum in the United States [5]. Offer preferences are also used in procurement of snow removal contracts [12] and have been studied in experimental settings [13]. A persistent finding is that procurement cost can be reduced by some positive preference, so long as the preference is not so strong as to inhibit participation by strong sellers.

Perhaps because most auction research in the area of offer discrimination considers singleunit auctions, set-asides and quotas, which apply only to multi-unit auctions, have received less study. In a multi-unit context quota neatly handles a problem inherent to the preference approach. To implement a price-preference mechanism, the auctioneer must have a good understanding of cost differences between classes. Such information is not necessary with quota. There are many applications, such as auctions for conservation land, when the auctioneer is less likely to know which classes are low-cost and which are high-cost, but nevertheless expects costs to differ widely across classes. In this case, providing a preference to the wrong party could increase procurement cost. Alternatively, the buyer can use a quota-a binding limit on the number of offers accepted from any one group-to encourage greater within-class competition, even if the buyer is unable to quantify cost differences across groups of sellers. Note that to enforce this rule, the buyer need not be able to identify which group of sellers is relatively low-cost. The only requirement is that the sellers themselves be aware of where costs in its own class lie relative to the broader population of classes.

Our research applies quota in an independent private values model of a one-shot, sealed-

bid auction. Because our focus is on procurement, we model a pay-as-offered auction as opposed to a uniform price auction. Almost all government procurement auctions use the pay-as-offered format. Thus our analysis differs from that of Ayres and Cramton [5] who investigate preferences and set-asides in multiple-round, open-bid auctions, and from Denes [6], who studies multiple auctions over time.

#### 2. The Environment and Symmetric Auctions

In this section we describe the environment of a generic procurement auction for multiple goods with two types of sellers. We then describe the methods we use to approximate equilibrium offer functions using strategies constrained to be flexible parametric functions of each type's cost.

Suppose two types of sellers, Type A and Type B, draw their costs randomly and independently from distributions  $F_A$  and  $F_B$  suppored by the intervals  $[\underline{c}_A, \overline{c}_A]$  and  $[\underline{c}_B, \overline{c}_B]$ . The distributions  $F_A$  and  $F_B$  differ by an additive parameter  $\delta$  such that  $F_A(c) = F_B(c+\delta)$ . This environment embodies a simple form of class asymmetry: sellers perceive Type A bidders as typically lower cost than Type B sellers but with identical within-group cost heterogeneity. There are  $N = n_A + n_B$  sellers, where  $n_A$  and  $n_B$  are drawn independently from  $F_A$  and  $F_B$ . The number of sellers of each type is common knowledge to the buyer and all sellers. The distributions are common knowledge to all sellers but not necessarily to the buyer. Each seller observes his or her own cost draw before submitting an offer but does not observe the cost draws of other sellers. Each seller submits a single offer and the buyer accepts m < Nof those offers.

We use this environment to investigate how offers, procurement costs and rents accruing to sellers change across different auction formats, both in theory and in the laboratory. First we consider theoretical predictions.

Because closed-form solutions to asymmetric auction problems typically do not exist, we solve for an approximate solution using numerical methods. To describe how these methods work, and verify that their accuracy, we begin with a symmetric auction ( $\delta = 0$ ) wherein the buyer accepts the lowest *m* offers, which has a well-known unique Bayesian Nash equilibrium:

**Theorem.** Let  $Y_1, Y_2, ..., Y_{N-1}$ , represent the random cost draws of each of the N-1 sellers that are not seller *i*, ordered from lowest-to-highest. Let *c* be seller *i*'s cost draw. The unique symmetric equilibrium strategy for all sellers when  $\delta = 0$  is  $b(c) = E[Y_m|Y_m > c]$ .

*Proof.* See Weber [14] or Ortega-Reichert [15].

In words, each seller offers his or her expectation of the  $m^{\text{th}}$  ordered cost draw conditional on that order statistic being greater than the seller's own cost draw (ie., conditional on winning). The intuition is as follows. A seller facing N - 1 ex ante identical competitors will submit an offer just low enough to be among the m lowest. Because the equilibrium is symmetric (every seller has identical beliefs about other sellers' costs), all players follow the same strategy and the offer function is monotonically increasing in c, so sellers with the lowest m cost draws will be accepted. Each seller thus forms expectations of what the  $m^{\text{th}}$ lowest cost draw will be, conditional on it being greater than the sellers own draw. Offers differ only by the extent to which sellers cost draws influence the conditional expectation. Expected procurement cost are thus  $mE[C_m]$  where  $C_m$  is the  $m^{\text{th}}$  lowest of N cost draws.<sup>4</sup> This result, though based on some technical and behavioral assumptions, gives predictions that turn out remarkably similar to experimental outcomes, as we show below.

The symmetric case gives the unique analytical solution described above. This solution does not hold for the asymmetric case where  $\delta > 0$ . Instead we develop numerical approximations for equilibrium bid functions. We use the symmetric case to verify the accuracy of the numerical approximation because it, unlike the asymmetric case, can be compared to the known equilibrium.

#### 3. Constrained Equilibria

Solving for equilibrium bid or offer functions when buyers or sellers are asymmetric is notoriously difficult [16]. We gain traction with this problem by using constrained strategies, a common approach that follows from Rothkopf's markup model [17]. Rothkopf considered strategies that were constrained to be a linear function of cost. Thus, each seller with a random opportunity cost c submits an offer  $\theta c$  that is some constant markup of cost, the parameter  $\theta$  having been chosen strategically prior to observing c to maximize expected gains. Linear strategies have two advantages. First, by constraining the strategy space one can solve for equilibria that would otherwise prove difficult or intractable. Second, strategies that are simple functions of a sellers private information can be more useful predictors of actual behavior than unconstrained Bayesian Nash strategies [18].

It is probably desirable, however, to consider constrained strategies that are more flexible than a linear markup. We therefore consider a family of strategies that can be made

<sup>&</sup>lt;sup>4</sup>Note that  $C_m$  is the  $m^{\text{th}}$  order statistic of N draws and thus differs from  $Y_m$  which is the  $m^{\text{th}}$  order statistic of N-1 draws.

arbitrarily flexible. Armantier and Richard [19] and Armantier, Florens, and Richard [2] show that strategies constrained to a sequence of progressively higher-order polynomials will converge to an unconstrained Bayesian Nash Equilibrium (BNE) should one exist. We can thus consider offer functions that are  $K^{\text{th}}$ -order polynomial functions of cost. Since offer functions differ across types, we denote the polynomial offer function for type T and order K by:

$$b_K^T = \sum_{k=0}^K \theta_k^T c^k.$$
(1)

We use a polynomial strategy function for simplicity, and because it is adequate in many cases. In general one can use any functional form. For consideration of asymmetric auctions in section 4 we use restricted cubic splines rather than polynomial. A constrained offer function leads to an associated equilibrium:

**Definition.** A constrained strategic equilibrium (CSE) is given by a pair of vectors  $\theta^* \equiv \{\theta^{A*}, \theta^{B*}\}$  and corresponding offer functions  $b_K^{T*}$  such that all sellers of type  $T \in \{A, B\}$  choose  $\theta^{T*} = \{\theta_1^{T*}, \theta_2^{T*}, ..., \theta_K^{T*}\} \in \Theta \subset \mathbb{R}^K$  to maximize his or her expected gain prior to observing cost and given all other sellers choose according to their constrained strategies defined by  $\theta^{T*}$ .

Unfortunately, it is not in general clear whether either a CSE or even an unconstrained equilibrium exists in an asymmetric auction. Moreover, existing convergence theorems say little about how far any finite order CSE may be from a true equilibrium.

One way to measure how well a CSE approximates a true equilibrium is to estimate the expected gains to a randomly chosen seller can earn by unilaterally deviating from the constrained strategy. That is, consider the expected gain to a randomly chosen seller *i* from choosing an offer unconstrained by the polynomial (or other finite order) offer function  $b_K^T$ , but assuming all sellers  $j \neq i$  choose according to the constrained strategy. If the CSE is close to a true equilibrium, the potential gains from unilateral deviation will be small.<sup>5</sup> This measure has strong intuitive appeal and earlier work on single-unit asymmetric auctions has shown high-ordered CSE have a low expected gain from optimal deviation [2]. This work also shows that for instances where there is a known equilibrium, high-order CSE equilibria

<sup>&</sup>lt;sup>5</sup>But there is no existing theory that implies that this necessarily goes the other way, that if the expected gains from optimal deviation are small then the constrained strategy is necessarily close to a true equilibrium. Thus, a mild incentive to deviate might only suggests a fairly plausible outcome.

well approximate that equilibrium.

**Definition.** The expected gain from unilateral deviation (EGD) from a constrained strategy defined by the finite parameter vector  $\theta$  is given by

$$EGD(\theta) = \mathbb{E}\left[ (b_i^* - c_i) \left( 1 - F_{j \neq i}(b_i^* | \theta) \right) - \left( b_K^T - c_i \right) \left( 1 - F_{j \neq i}(b_K^T | \theta) \right) \right]$$

where:

- $\mathbb{E}$  is the expectation taken over a randomly selected seller  $i \in \{1, ..., N\}$  with random and privately observed cost  $c_i$ ;
- $F_{j\neq i}(b|\theta)$  is the distribution function of the  $m^{th}$  ordered offer from sellers  $i \neq j$  all using constrained strategies defined by  $\theta$ , so that  $1 F_{j\neq i}(b|\theta)$  is the probability an offer of b is accepted;
- $b_i^*$  is the optimal unconstrained offer by individual *i* that maximizes  $(b_i c_i)(1 F_{i \neq i}(b_i | \theta))$ ; and
- $b_K^T$  is the  $K^{th}$ -order constrained strategy for individual i of type T.

Starting from this definition, we take an approach that is similar but distinctly different from CSE. The idea is simply that, if the size of the expected gain from unilateral deviation is a measure of closeness to true equilibrium, then a good choice for a constrained strategy that approximates the true equilibrium is one that minimizes the size of this expected gain.

**Definition.** The minimum expected gain from unilateral deviation (MEGD) is given by

$$MEGD(K) = \min_{\theta} EGD(\theta)$$

and the set of strategies selected to minimize this objective are defined as MEGD constrained strategies denoted  $\theta^*_{MEGD} \equiv \{\theta^{A*}_{MEGD}, \theta^{B*}_{MEGD}\}$  such that

$$\theta^*_{MEGD} \equiv \arg\min_{\theta} EGD(\theta) \tag{2}$$

In principal we can find  $\theta^*_{\text{MEGD}}$  for any constrained strategy of order K.

There are several advantages of this new approach. First, it directly optimizes the criteria by which closeness to true equilibrium has heretofore been judged. Second, the approach is much easier to program than CSE and uses simple computational methods, mainly because it transforms an equilibrium problem into a somewhat simpler optimization problem. Third, we found it much more stable than CSE and less sensitive to starting values. Fourth, for the symmetric auction in our environment (where  $\delta = 0$  or there is a single type) where there is a known and unique true equilibrium, we find this approach converges monotonically toward the true equilibrium as the order of polynomial increases. In contrast, CSE convergence is not always monotonic in practice.<sup>6</sup>

We illustrate MEGD strategies for a symmetric auction with 10 sellers, each with a privately observed cost drawn randomly from a uniform distribution over the range [0, 100]. The buyer accepts and pays as bid for the lowest six offers. This auction has a well-known solution given by the theorem above and is plotted by the solid black line shown the Figure 1. This true equilibrium offer function is derived from a series of conditional order statistics of cost draws. We found the function numerically by sampling nine random cost draws one million times, and finding average 6<sup>th</sup>-order statistic for the subset of cases where the order statistic was greater than each cost value along an index between zero and 100. This equilibrium offer function has an intuitive shape: At the low end the offer function begins at 60 and is flat, since low cost draws will almost surely will be among the lowest cost draws. As cost rises significantly, the offer function rises, gently at first, and then asymptotically approaching cost at the highest possible levels. This true equilibrium offer function serves as a benchmark for the constrained strategies we consider next.

The colored lines in Figure 1 show polynomial constrained strategies of order  $K = \{2, 3, 4, 5, 6, 7\}$ , each selected using MEGD. The coefficients for these functions are reported in table 1. These functions were found by numerically optimizing equation 2 using the algorithm outlined in the appendix.<sup>7</sup>. The table also reports a series of summary statistics for each MEGD constrained strategy: (1) EG, the expected gain to the seller (before observing cost); (2) EGD, the expected gain from optimally deviating from the constrained strategy after observing cost (the objective minimized); (3) EGD (%), which is EGD as a percent of EG; and (4) AD, the expected absolute deviation from the true equilibrium strategy.

<sup>&</sup>lt;sup>6</sup>This is true in our experience. Indeed, finding true CSE was difficult. Instead we found constrained strategies wherein the differential equations following from the first-order conditions were satisfied as closely as possible, but were not exactly zero. The CSE approach appears to be unstable because first-order conditions need only be satisfied in expectation, but may differ from zero by large amounts both positively and negatively for specific cost realizations, so long as positive deviations balance out with negative deviations. Put another way, it is theoretically possible to have a CSE with large EGD. In contrast, the MEGD approach penalizes both positive and negative deviations from zero across the whole range of possible costs.

<sup>&</sup>lt;sup>7</sup>The R computer code we used is also available upon request.

Figure 1: Polynomial constrained strategies that approximate true equilibrium strategies for a symmetric procurement auction



Notes: The graphs show polynomial constrained strategies for order two through seven for a symmetric auction in which privately observed costs are drawn randomly from a uniform distribution on [0, 100] and the buyer accepts the six lowest offers. The top graph shows the whole range of costs and offers and the bottom graph magnifies the lower, more critical range, where expected gains are largest due to high odds of offer acceptance and the large difference between the offer and cost. The true unconstrained theoretical equilibrium is given by the black line and the polynomial strategies are selected by <u>minimizing the expected gains from unilateral deviation from the constrained strategies (hence MEGD)</u>. At least in this case, as the order of polynomial increases, the MEGD constrained strategy converges monotonically toward the true equilibrium.

Both the graphs and the statistics show the MEGD polynomials approximate the true equilibrium very well. They also show strong convergence toward the true equilibrium strategy as the order of polynomial increases. The expected gain from optimally deviating from the quadratic strategy is just 0.0675 in comparison to an overall expected gain of 19.0, about one-third of one percent. Furthermore, in most cases, EGD declines by half or more with each successively higher polynomial. The expected gain from deviating from the seventh-order polynomial is less than 0.003, or 1.4 hundredths of a percent of the overall expected gain.

One notable feature of MEGD strategies is that they place more weight on fitting the lower portion of the offer curve than the higher portion. That is, the absolute difference between the polynomial strategy and the true equilibrium strategy tends to be smaller for lower-cost outcomes. This makes sense given expected gains are much greater for low-cost sellers as compared to high-cost sellers, which means the marginal gains from unilateral deviation are likely to be greater. High-cost sellers are unlikely to have their offers accepted anyway, and optimal offers tend to be much closer to cost, so marginal gains from deviation from the constrained strategy tend to be small. Thus, the MEGD approach chooses polynomials that work harder to fit the offer function on the lower end of the cost schedule. This is also why the absolute deviation (AD), while generally declining with higher-degree polynomials, does not decline monotonically. However, closer inspection of the lower portion of the offer function (the bottom panel of figure 1) shows strong convergence in the absolute sense for the lower part of the cost schedule. Even with magnification, for this portion of the curve the 7<sup>th</sup>-order polynomial is visibly indistinguishable from the true equilibrium.

#### 4. Approximate Equilibria for Asymmetric Auctions

Having shown that MEGD constrained strategies can well approximate true equilibria in a symmetric auction, we use the technique to find approximate equilibrium strategies for the asymmetric auctions that are the focus of our research. These auctions consider the environment described above with five Type A sellers and five Type B sellers and three cases of asymmetry,  $\delta = 25$ ,  $\delta = 50$  and  $\delta = 75$ . Thus, Type A sellers have costs drawn from a uniform distribution on [0, 100] and Type B sellers have costs drawn from a uniform distributions on [25, 125], [50, 150] or [75, 175], depending on the degree of asymmetry. There is common knowledge about the number of each type and the underlying distributions from which each type's costs are drawn, but information about the outcomes of the cost draws are the private information of each seller.

	Order of MEGD Polynomial					
	(2)	(3)	(4)	(5)	(6)	(7)
			Coefficient	ts		
Intercept	60.8	61	60.1	59.9	59.9	60
$\operatorname{Cost}$	-0.159	-0.061	-0.00115	0.0736	0.0462	-0.00593
$\mathrm{Cost}^2$	0.00545	0.00192	-0.00238	-0.00813	-0.00521	0.00195
$\mathrm{Cost}^3$		3.09e-05	0.000132	0.0003	0.000182	-0.000223
$\mathrm{Cost}^4$			-6.54e-07	-2.65e-06	-4.19e-07	1.08e-05
$\mathrm{Cost}^5$				7.85e-09	-1.18e-08	-1.75e-07
$\mathrm{Cost}^6$					6.53e-11	1.26e-09
$\mathrm{Cost}^7$						-3.51e-12
	L	Approxima	tion Statist	ics (see Not	es)	
$\mathbf{EG}$	19	19.4	19.2	19.2	19.2	19.2
EGD	0.0675	0.042	0.00776	0.00849	0.00543	0.00271
EGD (%)	0.355	0.216	0.0404	0.0442	0.0282	0.0141
AD	0.734	0.919	0.437	0.0408	0.059	0.128

Table 1: Summary of MEGD strategies for a symmetric auction

2

Notes: The table reports polynomial offer functions selected by minimizing the expected gains from optimally deviating from the constrained strategies (MEGD). The auction considered is a symmetric auction in which all sellers' costs are drawn from a uniform distribution on [0, 100] and the buyer accepts the six lowest offers. The offer functions are plotted in Figure 1 together with the true Bayesian Nash Equilibrium. The bottom four rows report other statistics of interest: EG is the expected gain to the seller before their random cost is observed; EGD is the expected gain from optimally deviating from the constrained strategy (the objective minimized); EGD(%) is  $100\% \times \frac{\text{EGD}}{\text{EG}}$ ; and AD is the expected absolute difference between the constrained strategy and the true equilibrium strategy.

A key difference between the asymmetric auctions and the symmetric auction above is that there are separate offer functions for each type. Thus, for a given order of approximation, there are twice as many parameters over which to minimize EGD. A second challenge with asymmetric auctions is that the odds that an offer from higher cost draws from the weak (higher cost) Type B sellers are extremely low, which makes these probabilities difficult to approximate from Monte Carlo simulation. The problem is accentuated with polynomial offer functions because offers near the boundary of the cost distribution can be very sensitive to coefficients on the higher-ordered terms. We therefore use *restricted cubic splines* instead of a simple polynomial. Restricted cubic splines, often used to fit non-linear statistical relationships, connect piecewise cubic polynomials across a series of knots, while constraining the function to be both continuously differentiable and linear outside the first and last knot. The linearity constraint near the boundary is well suited to asymmetric auctions because the equilibrium offer function is likely to be nearly linear at near the boundary anyway (since odds of acceptance are very nearly zero or very nearly one) and the linearity combined with second-order continuity at the knots makes the function robust to small simulation errors. We also found restricted cubic splines converged more quickly and had smaller MEGD than simple polynomials with the same degrees of freedom. Note that a restricted cubic spline requires two fewer degrees of freedom than the number of knots.<sup>8</sup>

For each degree of asymmetry, we compare three different auctions: (1) a uniform price auction that pays the lowest offer rejected (the seventh lowest of the ten) to those making the lowest six offers; (2) a pay-as-offered open auction in which the lowest six offers are accepted regardless of type; and (3) a pay-as-offered auction with quota in which the lowest six offers are accepted subject to the constraint that no more than four of five offers are accepted from any one type. In first auction all sellers offer their cost in equilibrium. In the second auction, sellers shade their bids to balance the odds of acceptance with the surplus received if accepted, but equivalence of the buyer's payment between (1) and (2) does not hold due to asymmetry. In the third auction imposes discrimination by instilling greater competition, particularly among lower-cost Type A sellers, because the type-specific limit is far more likely to bind for these sellers as opposed Type B sellers. We also consider a case wherein there is no true asymmetry but individual sellers are still assigned types. The idea here is to explore what happens if quota is imposed even when there is no true asymmetry.

The MEGD approximate equilibrium offer functions for both auctions are plotted in

<sup>&</sup>lt;sup>8</sup>As with the symmetric auction described above, details of the numerical implementation are given in the appendix and the computer code is available from the authors.

figure 2. The figure has four panels, one for each case of asymmetry, plus a "symmetric case" for which types are in fact identical are but are still potentially subject to type-specific quotas. Each panel plots the MEGD constrained strategies for the auction with quota and without quota, with each strategy constrained as a restricted cubic spline with six knots. A restricted cubic spline with six knots uses four degrees of freedom, plus one for the intercept, and thus uses the same number of parameters as a  $4^{th}$ -order polynomial. For all the constrained strategies plotted, the EG and EGD statistics (see table 2) are both absolutely and relatively less than that of the  $4^{th}$ -order polynomial constrained strategies for the symmetric auction plotted in figure 1, which were visibly indistinguishable from the known true equilibrium. The summary statistics were derived from simulating 10,000 auctions assuming sellers use the near-equilibrium MEGD strategies. The EGD statistics suggest the MEGD approximations are close to a true equilibrium, with an expected gain from unilateral deviation on the order of one-hundredth of one percent of the expected gain.

In addition to the EG and EGD statistics for each type of seller, table 2 reports the expected cost to the buyer for the six units purchased and the expected total cost to sellers of producing the six units purchased. Because the number of units purchased is fixed (and thus the benefit to the buyer), the lower the expected production cost of units purchased, the greater the total economic surplus.

Several interesting patterns emerge from the solutions to the asymmetric auctions. First, while auctions without quota are nearly as efficient as the uniform-price auction, the buyer's payment is less in the other two auctions. The non-equivalence of the buyer's payment between the pay-as-offered auction without quota and the uniform-price auction follows from asymmetry.

Second, in asymmetric auctions without quota, offers from the higher-cost Type B sellers are *less* than offers from the lower-cost Type A sellers when holding the specific cost draw fixed. (Focus on where the blue lines overlap.) This pattern is difficult to discern on the low-asymmetry case because the offer curves or so close, but very clear on the medium- and high-asymmetry cases. This result follows from the slightly different information sets that the two types possess: Type A sellers know they are competing against four Type A sellers and five Type B sellers, while Type B sellers know they are competing against five TypeA sellers and four Type B sellers. Thus, holding the seller's cost draw fixed, B types face greater competition than a A types, since B types face relatively more of the lower-cost Atypes. As Krishna [20] describes it, "weakness leads to aggression." This phenomenon leads

Figure 2: MEGD strategies for asymmetric auctions with and without quota



*Notes:* The graphs show MEGD constrained strategies for auctions with and without type-specific quotas. Each type-specific strategy is a natural spline with six knots.

to a very small amount of inefficiency.<sup>9</sup>

Third, in comparison to the non-discriminatory open auction, quota leads to lower offers from Type A sellers and higher offers from Type B sellers. Type A sellers make lower offers because while their offers could be among the six lowest, their offer may nevertheless rejected due to the quota. Conversely, Type B sellers face less competition from lower-cost A types, which implicitly increases demand for units from B types, who are much less likely to be rejected due to the quota. This effect is offset somewhat by the fact that Type A sellers are themselves bidding more aggressively, which feeds back and causes B types to bid more aggressively.

Fourth, if asymmetry is large enough, the quota mechanism saves procurement costs for the buyer in comparison to auctions without quota. In the medium-asymmetry case the buyer's expected payment is 528.7 with quota and 533.4 without; in the high-asymmetry case the buyer's expected payment is 581.5 with quota and 631.3 without. This result is not entirely obvious because the buyer will sometimes reject a lower-cost unit in favor of a higher-cost unit, but saves expenditures since low-cost sellers make more competitive offers. Except for the case with a small amount of asymmetry ( $\delta = 25$ ), savings from increased competitiveness outweighs losses from increased costs. Gains to the buyer from quota come at a cost to expected gains from *Type A* sellers and loss in total economic surplus, which is the standard result with imperfect price discrimination. Not surprisingly, cost savings to the buyer from using quota are larger the greater the amount of asymmetry.

Fifth, in the low asymmetry case quota barely reduces the buyer's expected payment, the increase is tiny (457.01 with quota and 457.06 without quota). Even in the case of no asymmetry, expected buyer costs are are just 0.76, or less than 2 *tenths* of one percent, greater than an auction without quota. This is interesting and practically useful because in real-world applications buyers may have less information about the degree of asymmetry across types than sellers do. If such a seller were to use the quota mechanism, there is little additional payment in the event asymmetry is actually too low for profitable use of quota, but large savings in the event asymmetry is high. Intuitively, the lower the asymmetry, the less likely the quota will bind in the first place. Consider, for example, the case both types were to have their costs drawn from the same distribution ( $\delta = 0$ ) and thus have identical offer functions. The probability that the quota binds for any one type would therefore equal the odds that the largest of five uniform cost draws is smaller than the second smallest

<sup>&</sup>lt;sup>9</sup>In the simulations, efficiency, indicated by the sellers' total cost, is sometimes slightly higher in the pay-as-offered auctions without quota, but this only happens due to chance error in the simulation.

	Buyer's	Type A	Type B	Type A	Type B	Sellers'
	Payment	EG	EG	EGD	EGD	Total Cost
			No Asymi	metry, $\delta =$	0	
Uniform Price	380.76	18.99	19.09	0	0	190.35
No Quota	381.53	19.15	19.17	0.01	0.01	189.96
Quota	382.29	19.18	19.19	0.00	0.01	190.46
		Ι	low Asym	metry, $\delta =$	25	
Uniform Price	457.72	28.48	11.64	0	0	257.12
No Quota	457.06	27.70	12.40	0.01	0.02	256.55
Quota	457.01	26.93	12.80	0.02	0.02	258.37
		Me	dium Asy	mmetry, $\delta$	= 50	
Uniform Price	540.89	40.62	6.30	0	0	306.33
No Quota	533.13	38.14	7.25	0.03	0.04	306.17
Quota	528.15	32.70	10.21	0.02	0.04	313.61
		H	ligh Asym	metry, $\delta =$	= 75	
Uniform Price	655.31	59.21	3.66	0	0	340.96
No Quota	631.17	53.38	4.59	0.04	0.05	341.29
Quota	581.34	33.66	9.75	0.01	0.04	364.27

Table 2: Summary of MEGD outcomes for asymmetric auctions

*Notes:* The table reports average outcomes from simulating 10,000 auctions using MEGD constrained strategies plotted in figure 2. Type A sellers' costs are drawn randomly and independently from a uniform distribution on [0, 100]; Type B sellers' costs are drawn from a uniform distribution on  $[0+\delta, 100+\delta]$ . There are five sellers of each type and the buyer (auctioneer) purchases six units in each auction. In Uniform Price auctions, the lowest six offers are accepted, each is paid the lowest rejected offer (the seventh ordered offer), and all offers equal cost and the auction is efficient. In No Quota auctions the lowest six offers are accepted regardless of type. In Quota auctions the lowest six offers are accepted subject to the constraint that no more than four of five offers from any one type are accepted. Buyer's Payment is the average sum of payments by the buyer to the sellers for the six units purchased.  $EG^A$  and  $EG^B$  are the expected gains to each type of seller.  $EGD^A$  and  $EGD^B$  are the expected gains from unilateral deviation, a measure for how close the MEGD strategies are from true equilibrium (zero would be a true equilibrium). Sellers' Total Cost is the expected cost to sellers of the six units purchased by the buyer. Holding the buyer's value of the six units purchased fixed, a lower cost implies greater total economic surplus.

of two uniform draws, which equals 0.024. The slightly lower odds of selection causes the equilibrium offer functions to be only slightly higher than they would be in an open auction without quota. An inefficient selection of offers occurring with a probability of slightly less than five percent, would create a second small inefficiency and cost to the buyer. But this is the worst possible scenario from a naive application of a quota.

#### 5. Experiments

There are several reasons for experimentally testing the quota mechanism in the laboratory. First, because closed-form theoretical predictions are difficult or impossible to derive, empirical validation of the predicted influence of quota complements the approximate theoretical solutions already described. Second, behavior may systematically differ from theory. Human sellers do not always use strategies consistent with standard game theoretic concepts, even in environments where the mathematical theory suggests such optimal strategies are simple, and in this case the strategies are not particularly simple.<sup>10</sup> Third, we wish to test whether the price discriminating mechanism is implementable in the sense that the rules can be easily explained to sellers and outcomes are reasonably robust to tacit collusion or to poor, non-strategic offer making by some sellers.

#### 5.1. Experimental Procedure

We report results from 17 experimental sessions that included 269 individual real-money auctions, plus demonstration auctions. Each session included ten undergraduate students from the University of Maryland. All experiments were computerized, using custom-made software. In each session, five participants were randomly classified as Type A sellers and five were classified as Type B sellers. Subjects did not know the Type-identities of their competitors, but did know that there were five sellers of each type. Subjects were also told how their costs were drawn randomly, with Type A costs on the interval from \$0.00 to \$100.00 (rounded to the penny) with all values equally likely to occur, while Type B sellers costs were drawn uniformly from the interval \$50.00 to \$150.00. This environment is equivalent to the medium asymmetry case examined numerically above.

Subjects were randomly assigned to a computer terminal that displayed information for them when they sat down. They were shown written instructions that were simultaneously read aloud by the proctor, and were then given time to re-read the instructions on their own

<sup>&</sup>lt;sup>10</sup>It is well known, for example, that bidders in second-price sealed-bid auctions typically do not submit bids equal to their costs, even though doing so is unambiguously optimal [21]

before the first auction began.<sup>11</sup> Each subject had an opportunity to practice each auction hypothetically before participating for real money. Payments and costs were delineated in "dollars" for simplicity, but an exchange rate between experimental dollars and U.S. dollars made experimental dollars was worth much less than U.S. dollars.<sup>12</sup> Students earned an average of about \$30 in an hour long session.

Each subject participated in both open auctions and auctions with type-specific quotas. Because every subject participated in both treatments, we can make both within-subject and between-subject comparisons. We varied the order of treatments to control for learning effects and, in some sessions, reverted to the first auction after several rounds of each treatment (with or without quota). Subjects did not change between Type A and Type B in sequential auctions, mainly to facilitate learning of optimal type-specific strategies.<sup>13</sup> Subjects would, however, receive new random cost draws in each auction. These variations allowed us to examine whether individual offer behavior varied with experience and to control for these effects if necessary. In the regression analysis, we explicitly model order and session effects using random and fixed effects. These have little influence on estimated treatment effects but do influence standard errors. Perhaps most importantly, we find no evidence of last-round effects.<sup>14</sup>

We focus on the medium-asymmetry case in our experimental analysis. While it may be worthwhile to confirm that procurement cost is further reduced by a quota as asymmetry increases, we were concerned about perceived fairness in the experimental auctions. With large asymmetry, high-cost Type B sellers rarely have their offers accepted; and when they are accepted, they tend receive very small surplus. The small rewards for participation can

<sup>&</sup>lt;sup>11</sup>Instructions and pictures of the computer screen are provided in an appendix.

<sup>&</sup>lt;sup>12</sup>One experimental dollar was worth \$0.05. We also gave a fixed dollar amount for participating, typically \$5. Average earnings by subjects during the auction session were about \$30, but varied substantially depending on whether the subject was a Type A or a Type B bidder (standard deviation of earnings was approximately \$19). To even-out earnings we allowed subjects to play a very short game at the conclusion of the session in which the subjects that had been disadvantaged during the auction enjoyed more favorable terms.

<sup>&</sup>lt;sup>13</sup>We did, however, experiment with switching types across treatments.

<sup>&</sup>lt;sup>14</sup>Because each group of students played multiple rounds of each auction, some may argue that our experiment is truly a dynamic game, not the static one about which we have theorized. By this reasoning, the final rounds may be the only true experimental trials of our theoretical environment. We recognized this point but generally believe dynamic strategic considerations to be of relatively small importance, particularly in an environment with a relatively large number of players making anonymous offers. Such an environment makes it particularly difficult to instill rewards or punishments that can sometimes make repeated play differ from single-shot games. In any case, a Nash equilibrium in each stage-game of repeated play is itself a Nash equilibrium, as well as a subgame-perfect Nash equilibrium, so our static predictions nevertheless give one plausible solution to the dynamic game.

lead to frustration, loss of interest and unrealistic behavior. It seems intuitively clear that the benefit of using a discriminating mechanism should increase as seller heterogeneity increases. Consequently we didnt believe that it was necessary to test the quota auctions under the high asymmetry condition. We did not test the quota mechanism with low between-group heterogeneity, mainly because with low asymmetry, theory suggests quotas will have little influence on offer strategies or outcomes. However, given that we observe more aggressive bidding in our experimental sessions than we expected, in retrospect this may have been desirable. We leave experimental evaluation of quota in low asymmetry cases for future research.

#### 5.2. Experimental Data

In figure 3 we present experimental offers plotted against cost draws. The top two panel of the figure show offers from auctions without quotas and the bottom two panels show offers from auctions with type-specific quotas. Raw individual offers are presented in scatter plots on the left and smooth cubic polynomials summarizing the offer trends are shown on the right. The different colors, point types and line types indicate different experimental sessions, each of which included ten or more individual auctions. Statistics summarizing offers broken out by seller type (A or B) and cost quartile are reported in table 3. While individual offers vary considerably conditional on cost (a phenomenon contrary to theory but not altogether surprising) offers do cluster around the same regression line conditional on cost.<sup>15</sup>

In auctions without quota, offers increase slightly with cost on the lower half of the cost schedule and are centered around the average highest offer accepted, about \$85. When cost draws are above the average marginal offer, offers tend to increase more sharply with cost, with a slope equal to about 0.8. Interestingly, offers in the first and second quartiles of Type B sellers are somewhat less that those in the third and fourth quartiles of Type A sellers, which have similar cost draws. This result is consistent with the theoretical prediction that "weakness leads to aggression." These differences are statistically significant and similar in magnitude to those predicted by the approximate theoretical equilibria reported above.

In auctions with quota as compared to auctions without quota, we see significantly lower offers from Type A sellers that tend to increase somewhat more sharply with cost. Offers from

<sup>&</sup>lt;sup>15</sup>Some individual offers appear to make little sense. For example, some subjects submitted offers less than cost and others submitted offers in excess of 300, but these are rare and thus are not included in the analysis. A total of 37 of 1950 offers were dropped from auctions without quota (21 below cost, 16 above 300) and a total 17 of 740 offers dropped from auctions with quota (13 below cost, 4 above 300).



Figure 3: Individual Offers from Experiments

*Notes:* The top graphs show individual offers from experimental auctions plotted against sellers' privately observed costs together with fitted cubic polynomials of offer against cost. All experiments included 10 sellers, five of each type, with cost draws uniformly drawn from [0, 100] and [50, 150] respectively. The top tow panels show offers from experiments without quota and the bottom two panels show offers from experiments without quota and the bottom two panels show offers from experiments with quota. Different colors, point and line styles delineate different experimental sessions (i.e., participants are held fixed for a given color and point style). Offers below cost (21 without quota and 4 with quota) and offers above 200 are not shown.

		Type A	Sellers			Type B	8 Sellers	
Cost Quartile:	$1^{st}$	$2^{nd}$	$3^{rd}$	$4^{th}$	$1^{st}$	$2^{nd}$	$3^{rd}$	$4^{th}$
				Without	Quota ( $\delta = 50$	)		
Mean Cost	13.5	38.5	63.2	87.7	62.0	86.4	110.3	135.5
(SD Cost)	(7.9)	(6.7)	(7.9)	(6.8)	(7.9)	(6.7)	(7.9)	(6.8)
Mean Offer	81.6	83.6	90.3	99.8	$86.7^{\#}$	$95.7^{\#}$	118.3	144.0
(SD Offer)	(21.3)	(14.5)	(16.7)	(16.2)	(14.6)	(10.4)	(16.8)	(15.5)
Ν	242	241	241	242	237	237	236	237
				With Q	uota ( $\delta = 50$ )			
Mean Cost	12.2	38.0	64.7	87.6	62.9	86.0	111.2	135.4
(SD Cost)	(6.0)	(8.5)	(6.8)	(7.5)	(6.0)	(8.5)	(6.8)	(7.5)
Mean Offer	$64.5^{*}$	$71.4^{*}$	$81.6^{*}$	$93.8^{*}$	82.7*	95.5	115.5	140.5
(SD Cost)	(16.2)	(14.4)	(10.2)	(11.5)	(9.7)	(12.6)	(8.2)	(13.2)
Ν	92	92	91	92	89	89	89	89

Table 3: Summary of Experimental Offers

Notes: The table summarizes experimental outcomes from asymmetric procurement auctions with ten sellers, five of Type A and five of Type B, each with one unit to sell. Type A sellers' costs are drawn randomly and independently from a uniform distribution on [0, 100]; Type B sellers' costs are drawn from a uniform distribution on [50, 150]. The buyer (auctioneer) purchases six units in each auction. In auctions without quota, the lowest six offers are accepted regardless of type. In auctions with quota, the lowest six offers are accepted subject to the constraint that no more than four of five offers from any one type are accepted. A (\*) indicates quota offers that are significantly less than offers without quota in the same cost quartile, with a p-value less than 0.005. In auctions without quota, a (#) indicates Type B offers that are significantly less than Type A offers with the same distribution of costs, again with a p-value less than 0.005. Seemingly ill considered offers—those less than cost or greater than 300—were dropped from the sample. A total of 37 of 1950 offers were dropped from auctions without quota (13 below cost, 4 above 300).

Type A sellers with costs below 50 are 10 to 17 less in auctions with a quota as compared to auctions without quota; offers from Type A sellers with costs between 50 and 100 are 7 to 9 less with quota as compared to auctions without quota. Thus, the prediction that type-specific quotas increase competition among low-cost types and that this causes them to submit more aggressive offers is firmly realized in the laboratory.

Experimental offers from Type B sellers in auctions with quota display the clearest departure from theory, especially offers from cost draws in the first quartile (approximately [50,75]). In theory, these offers are higher than offers in auctions without quotas, because the quota on Type A sellers artificially increases demand from B types. In experiments, however, offers from B types were lower in auctions with quotas than without quotas. Indeed, in auctions with quota, Type B sellers submitted offers only slightly higher than those from A types with similar cost draws, and not significantly so. Thus, in experiments, the "weakness leads to aggression" phenomenon is just offset by increased demand from weaker B types, where in theory greater demand from B types has a much larger influence on offers. This phenomenon makes use of quota relatively more attractive to the buyer in practice than it is in theory. This phenomenon observed for relatively low-cost B may also push A types toward somewhat lower offers, as Type A sellers do make offers that are generally below those predicted by theory.

Where the buyer gains from low-cost types submitting offers somewhat closer to their costs, there are two tradeoffs associated with using quota to incite these lower offers. First, high-cost sellers bid less aggressively, because the quota on low-cost types artificially increases demand from high-cost sellers. In the experiments, however, offers from higher-cost Type B sellers are still slightly lower than those without a quota. Second, and more importantly, ex post the buyer potentially sacrifices a lower offer from an A type in favor of a higher offer from a B type in order to enforce the quota. Despite these tradeoffs, the gains to the buyer from lower offers exceeds the tradeoffs, with quota reducing the cost of procurement by 47.5 or 9.1% in comparison to auctions with Quota. In theory, the savings from quota conditional on the same cost draws would have been just 2.5 (less than 0.5%). These statistics, reported in table 5.2, do not account for order effects or variation across individuals or experimental sessions.

It is also interesting to compare the social efficiency of auctions with and without quota, both in the laboratory and in theory. These comparisons can be made by looking at the average total cost of six sellers, reported in the last row of table 5.2. Because six units are purchased in all auctions, thereby holding fixed the benefits of goods produced and

	Theoretical	Outcomes <sup>a</sup>	Experimenta	al Outcomes
	No Quota	Quota	No Quota	Quota
Buyer's Payment	533.7	531.3	520.0	472.5
Average Gain, Type A	37.2	32.0	32.7	23.5
Average Gain, Type B	7.3	9.7	5.9	4.9
Sellers' Total Cost	311.0	323.2	326.9	330.4

Table 4: Basic Comparison of Theoretical and Experimental Auction Outcomes

 $^{a}$  Here the theoretical outcomes are matched to the specific cost draws that came up during the experiments and thus differ slightly from the values reported in Table 2 due to chance variation in cost draws.

bought, social surplus is maximized when sellers' total costs are minimized. In theory, auctions without quota are the most socially efficient, although not *perfectly* efficient due to the weakness-leads-to-aggression phenomenon, which makes offers slightly non-monotonic in cost. In theory, the imposition of quota raises sellers' total cost by 12.2 (323.2 versus 311.0), even though the buyer gains slightly. The buyer's gains are realized via lower surplus to Type A sellers (32.0 versus 37.2 per seller) that are partially offset by higher gains to Type B sellers (9.7 versus 7.3). In experiments, however, the buyer's much more substantial gain from quota (an average payment of 472.5 versus 520.0) comes at a lower social cost of just 3.5 (330.4 minus 326.9). In experiments, quota lowers Type A surplus more substantially than it does in theory and slightly lowers Type B surplus as well.

While the difference in social costs between auctions with and without quota is smaller in experiments than in theory, due to somewhat erratic offers in the laboratory (ie., apparent random variation conditional on cost and type), social costs are somewhat higher overall in the laboratory as compared to theory.

#### 5.3. Regression Models of Sellers' Offers

We use regression analysis to estimate offer tendencies more precisely and compare them statistically across quota and no-quota auctions, and against theoretical offers derived above. All specifications have the form:

$$Offer_{sit} = \alpha + \beta O_{sit} + I(Quota \times Type)_{sit} \Gamma + I(Quota \times Type)_{sit} \times O_{sit} \Theta + controls + \epsilon_{sit}$$
(3)

where O denotes the MEGD theoretical offer associated with the seller's cost draw, type, and whether or not a quota was imposed;  $I(\text{Quota} \times \text{Type})$  is vector of (0-1) indicator variables delineating seller type and whether or not an auction was imposed, and  $\alpha$ ,  $\beta$ ,  $\Gamma$ , and  $\Theta$ are parameters. The subscripts *s*, *i* and *t* denote, respectively, the experimental session, individual seller, and round number within a session. Under the null hypothesis that theory is exactly correct,  $\alpha = 0$ ,  $\beta = 1$  and all other coefficients, including the error variance, equal zero. The coefficient vectors  $\Gamma$  and  $\Theta$  allow departures from theory to vary across treatments (quota or no quota) and for different seller types. In all specifications, the indicator variable left out is for *Type A* sellers in auctions without quota, so  $\alpha$  and  $\beta$  reflect this group. The controls vary by specification and are used to account for round, individual and experimental session effects. Some specifications account for these factors using fixed effects; others use random effects.

A summary of results from offer-level regressions reported in table 5.3. The first specification (column 1) is a baseline model that includes only an a round effect. The round variable indicates the auction number within an experimental session and treatment. The second specification (column 2) also includes fixed effects for each experimental session. The third specification (column 3) interacts session effects and round effects, allowing each session to have its own intercept and trend across repeated auctions. The fourth specification models session-by-round effects as random effects rather than fixed effects, and allows these effects to be correlated. The fifth and last specification (column 5) models individual seller-by-round random effects. The random effects specifications may be more appropriate if attempting to draw inferences about offer patterns for subjects not included in the experiment. Such an interpretation requires an assumption that our subjects were drawn randomly from a larger population, which was not the case. Subject participation was voluntary, and thus unlikely representative, but we see no reason why selection might be biased toward a specific kind of strategic behavior. While not testable, we find it plausible that selection effects are statistically ignorable. We restricted the sample to offers from non-demo rounds that were greater than cost and less than \$300. Out of 2690 offers, this excludes 34 offers that were less than cost and 20 offers that were greater than \$300. Except for the  $R^2$  values, these omissions have little influence on the results.

Although the data reject the theory in the literal sense, offers conditional on cost are centered very closely around MEGD equilibria. This result is most true for *Type A* sellers in auctions without quota. For this group, we fail to reject the null hypothesis that the mean offer conditional on cost equals the theoretical offer ( $\alpha = 0$  and  $\beta = 1$ ). We can see this graphically in the top-left panel of figure 5.3. The fitted regression line is indistinguishable from the line where offer equals the theoretical offer.<sup>16</sup> Still, while the theory seems to provide some useful guidance about offer tendencies, variation around the theoretical offer is not predicted by Bayesian Nash equilibrium, and so literal acceptance of the theory is strongly rejected.

All specifications show somewhat larger differences from the theoretical predictions in auctions with quota. Offers with quota generally have a lower intercept and increase more than one-for-one with theoretical offers. Thus, offers with lower cost draws tend to be below that predicted by theory and offers with higher cost draws tend to be above that predicted by theory. This pattern can be seen both in the coefficient estimates and from inspection of the lower two panels in figure 5.3, and holds up after taking into account round and session effects. The net effect of quota is to lower the overall costs of procurement relative to theory, because lower offers (and lower cost draws) are more likely to be accepted.

Figures 5 and 6 with show the relationship between theory and experimental offers for each experimental session, after removing individual intercept and round random effects. These are predictions from the model reported in column 5 of table 5.3. Theory predicts that all offers would lie exactly along the diagonal of each panel in the figure.

Interestingly, while the round effect is both small and statistically insignificant by itself, it is strongly significant when interacted with session or individual effects. When modeled as a random effect, the standard deviation of the round effect, whether with outcomes that vary across individuals or experimental sessions, is also a large and statistically significant. What the data show is that, while offers often begin well above or below theoretical predictions, they generally trend toward equilibrium predictions or a little below in successive rounds. In quota sessions, however, offers are more likely to start low and trend up toward theoretical predictions, while auctions without quota often started both above and below equilibrium. We illustrate these effects graphically in the next section where we consider auction-level regressions.

#### 5.4. Regression Models of Buyer's Total Payment

Here we examine the buyer's total payment to sellers for the six units purchased. The unit of analysis is therefore an auction, not an individual offer. For each experimental auction we found the theoretical offer associated with each cost draw and type, selected the theoretically winning offers, and calculated the total payment the buyer would have paid had all offers

<sup>&</sup>lt;sup>16</sup>Note that the bunch of points lined up near the vertical axis follows from the fact that the theoretical offer curve is generally flat over its lower portion.

	OLS Models Mixed 1			Mixed Eff	fects Models
	(1)	(2)	(3)	(4)	(5)
		Fixed	Effects	Estimate/(Standa	ard Error)
$Intercept^a$	-10.22	-12.40	-10.62	-12.69	-10.09
	(6.54)	(5.83)	(6.03)	(6.30)	(5.41)
$\operatorname{Round}^a$	-0.11	-0.09	-0.16	-0.06	-0.10
	(0.06)	(0.05)	(0.24)	(0.13)	(0.07)
Theoretical Offer	1.10	1.13	1.11	1.12	1.11
	(0.07)	(0.06)	(0.06)	(0.06)	(0.06)
No Quota & Type B	7.99	12.19	12.24	12.41	9.50
	(6.97)	(6.13)	(6.09)	(6.09)	(5.77)
Quota & Type A	-76.9	-66.6	-68.3	-68.0	-63.9
	(12.8)	(11.3)	(11.2)	(11.6)	(10.4)
Quota & Type B	-23.2	-17.2	-18.2	-19.5	-22.4
	(8.2)	(7.2)	(7.2)	(7.8)	(6.8)
$Theory \times$	-0.07	-0.12	-0.11	-0.12	-0.10
No Q & Type B	(0.08)	(0.07)	(0.07)	(0.07)	(0.06)
$Theory \times$	0.86	0.79	0.81	0.79	0.72
Quota & Type A	(0.15)	(0.13)	(0.13)	(0.13)	(0.12)
$Theory \times$	0.15	0.13	0.14	0.13	0.14
Quota &Type B	(0.08)	(0.07)	(0.07)	(0.07)	(0.07)
Session Fixed Effects	No	Yes	No	No	No
Session $\times$ Round Effects	No	No	Yes	No	No
Session $\times \ln(\text{Round})$ Effects	No	No	No	No	No
	Ran	ndom E	ffects Es	stimate/ $(95\%$ Con	fidence Interval)
SD Intercept				10.7	13.7
Group = Session in (4), In	ndividua	l in (5)		(7.45, 15.3)	(12.1, 15.6)
SD of Round				0.47	0.61
				(0.28, 0.81)	(0.46, 0.81)
$\operatorname{Corr}(\operatorname{Int}, \operatorname{Round})$				-0.89	-0.92
				(-0.97, -0.65)	(-0.97, -0.79)
Adjusted $\mathbb{R}^2$	0.682	0.755	0.760		
RMSE	14.56	12.74	12.58	12.57	10.71

Table 5: Summary of Offer-Level Regressions

 $^{a}$  For models with session or individual fixed effects or random effects, the coefficient and standard error reported for the intercept and/or round is for mean of the effect across all individuals or sessions.



Figure 4: Comparison of Experimental and Theoretical Offers by Treatment and Type

Notes: The figure shows experimental offers in relation to theoretical (MEGD) offers by treatment (Quota or No Quota) and type of seller (A or B). The blue lines show fitted regression lines; the red dashed lines are where Theory = Actual.



Figure 5: Offers Versus Theory by Session Without Random Effects: No Quota

*Notes:* The figure shows theoretical offer predictions in relation to offers from actual experiments with session and round random effects removed. Each panel shows offers from an individual experimental session *without* quota. A perfect diagonal line indicates where actual offers equal theoretical offers.



Figure 6: Offers Versus Theory by Session Without Random Effects: With Quota

*Notes:* The figure shows theoretical offer predictions in relation to offers from actual experiments with session and round random effects removed. Each panel shows offers from an individual experimental session *with* quota. A perfect diagonal line indicates where actual offers equal theoretical offers.

accorded exactly with theory. Our analysis focuses primarily around the difference between the actual amount the buyer paid in relation to the theoretical payment. A summary of these basic results is presented in figure 7, which shows actual total payments in relation to theoretical payments, with the kind of auction (with or without quota) indicated by point color. Note that the only factors driving variation in the theoretical total payment are random variation in cost draws and whether or not the auction has type-specific quotas.

The models, summarized in table 6, regress the actual total payment against the theoretical payment and auction characteristics. All models have the form:

$$Payment_{st} = \alpha + \beta P_{st} + \gamma I(Quota)_{st} + \theta I(Quota)_{st} P_{st} + controls + \epsilon_{st}$$
(4)

The specifications are similar to the offer-level regressions, except without type effects, because both types are represented in every auction. The indicator variable is therefore a scalar for whether or not a quota is imposed, not a vector crossing both auction and seller types. The variable P represents the theoretical payment—that predicted by MEGD strategies conditional on cost draws for the particular auction, with or without quota. The subscripts sand t denote the session and round, respectively, and controls include different specifications of round and session effects. We also report regressions that omit the theoretical payment and instead substitue the average cost draws for each seller type. These regressions provide direct evidence of the overall empirical effect of quota on procurement cost, not just the difference between the theoretical and empirical effect of quota.

Results are reported in table 6. Column 1 reports a basic regression of the total payment against P, the quota indicator, round, and session fixed effects. Column 2 adds session fixed effects. Column 3 adds session-specific round effects. Column 4 models session and session-specific round effects as random rather than fixed. Columns 5 and 6 are like models reported in columns 3 and 4 except the theoretical payment is replaced with controls for average Type A and Type B cost draws.

The importance of the interaction effect between round and session effects is shown by the increase in  $\mathbb{R}^2$  in columns 2 and 3 and strong statistical significance of the random round effects in columns 4 and 6. It is also illustrated by the sharp reduction in RMSE. We illustrate these effect in figure 8. The figure plots the different between the experimental payment and the actual payment in relation to round number. When we connect individual sessions with line segments, it is clear how the buyer's total payment converges toward the theoretically predicted payment, or a bit below, in later rounds.

Because offers are highly correlated across experimental sessions and not all sessions





*Notes:* The figure shows the buyer's total payment to the six winning sellers in relation to the MEGD derived theoretical payment. Auctions without quota are in blue and auctions with quota are shown in black. Simple regression lines are also drawn for each kind of auction. The red dashed line shows indicates where actual payment equals the theoretical payment.

include both quota and no-quota treatments, it is important to consider whether session effects confound the overall effect of quota or make quota effects more uncertain. The random effects models (columns 4 and 6) allow us to draw some inferences about the effect of quota on the population of possible sessions that might have occurred.

After accounting for session effects, we generally find similar results regardless of specification. While the buyer's savings from quota are significantly less than predicted by theory, they are not as large as suggested by the summary statistics, due to the large influence of session effects. In theory, quota saves about 5 experimental dollars, or about one percent of the cost from an auction without quota. The empirically estimated savings are about 13. The additional savings generally come from lowest cost sellers making offers lower than predicted by theory. This is indicated by the negative coefficient on the quota indicator variable and and positive value on quota interaction terms.

One indication for the usefulness of the theoretical predictions comes from comparing columns 3 and 5 with columns 4 and 6, respectively. These pairs of models are similar except that, in the latter two cases, cost draws for each type are substituted for theoretical predictions used in the first two cases. These comparisons show that the two models using theoretical predictions fit the data better than models using raw cost draws, even though the models with raw cost draws use more degrees of freedom.

#### 6. Conclusions

This paper studies a new kind of auction for asymmetric, multiple-unit procurement auctions and compares this auction, both theoretically and experimentally, to a standard pay-as-offered multiple-unit procurement auction. The new auction simply imposes a quota on the number or share of offers accepted from each identifiable class or type of sellers. The quotas implement a discriminating mechanism that will generally allow a buyer to reduce procurement costs in a manner that mimics bid preferences, which have been studied in single-unit asymmetric auction settings. A key advantage to quotas over other forms of type discrimination is that the buyer can save costs without knowing the underlying cost distributions of different types, which types have lower costs, or whether there are even significant differences in cost distributions across types. If observed types are actually symmetric then there is almost no loss of efficiency or of procurement cost from using quotas; conversely, if asymmetries are large, the buyer's savings from quota could be substantial.

In the laboratory, we find offers appear roughly consistent with theory but differ in two key respects. First, unsurprisingly, individual offers vary randomly around the theoretical

			Regress	sion Mode	el	
Explanatory	(De	pendent	Variable:	Total Pay	ment to Se	ellers)
Variable	(1)	(2)	(3)	(4)	(5)	(6)
		-	Estimate/(	Standard E	rror)	
Intercent	-92.1	-163	-23.5	426	-46.2	417
mercept	(149)	(68.1)	(287)	(68.1)	(81.9)	(28.7)
Round Number	0.869	0.534	0.402	0.521	0.461	0.59
Round Number	(0.711)	(0.416)	(1.69)	(1.74)	(1.35)	(1.4)
Theoretical Payment	1.13	1.27	1.01		1.06	
Theoretical Tayment	(0.278)	(0.16)	(0.15)		(0.148)	
Quota	-360	-234	-386	-183	-366	-179
Quota	(235)	(134)	(124)	(36.2)	(124)	(36.2)
Quota Theoretical Payment	0.593	0.427	0.709		0.67	
Quota× i neoreticar i ayment	(0.442)	(0.252)	(0.234)		(0.233)	
Avorago Cost Type A Sollors				0.983		0.954
Average Cost Type A Beners				(0.170)		(0.170)
Avorago Cost Typo B Sollors				0.412		0.513
Average Cost Type D Sellers				(0.174)		(0.171)
$Ouota \times AC$ Type A				0.661		0.715
Quota×AC Type A				(0.31)		(0.31)
$Ouota \times AC$ Type B				1.38		1.3
Quota×AC Type D				(0.311)		(0.31)
Adjusted R-squared	0.225	0.762	0.810	0.799		
RMSE	56.2	30.4	26.6	27.2	26.7	27.4
Session Fixed Effects	No	Fixed	Fixed	Fixed	Random	Random
Session-Specific Round Effects	No	No	Fixed	Fixed	Random	Random
Theoretical Savings	4.08	4.08	4.08	4.08	4.08	4.08
from Quota	4.50	4.50	4.50	4.50	4.50	4.50
Estimated Savings	<i>4</i> 8 70	11 72	12.06	12 20	13/13	13 11
from Quota	40.19	11.14	12.90	12.29	10.40	10.11

Notes: The table reports regressions of the buyer's payment (the sum total paid to the six winning sellers) against the theoretical payment and other explanatory variables. If sellers behaved perfectly according to theory, the coefficient on Theoretical Payment would be 1, all other coefficients would be zero, and the  $\mathbb{R}^2$  would be 1. Thus, regressions in columns (1) through (3) estimate the effect of quota on buyer's payment after accounting for the theoretical influence of quota. Regressions in columns (4) through (6) estimate the overall influence of quota on buyer's payment. In models with session effects, the reported intercept and/or round effect is for the estimated mean across sessions. Estimated savings from quota include both theoretical and estimated empirical effects of quota. The RMSE is the standard deviation of residuals or, in models with random effects, the standard deviation of residuals conditional on random effect outcomes.

Figure 8: Buyer's Experimental Payment Less Theoretical Payment by Auction Round



*Notes:* The figure shows the difference between the buyer's actual payment and the MEGD theoretical payment in relation to the auction round. Each line connects auctions from a single experimental session with a fixed group of sellers.

offer conditional on cost; second, offers with quota tend to be somewhat less than predicted by theory, especially for low-cost sellers, thereby yielding greater savings for the buyer than predicted by theory. While buyer savings come at a cost to social efficiency and gains to sellers, efficiency losses from quota are also smaller in the laboratory than they are in theory.

A key challenge to conducting this study was deriving clear theoretical predictions to compare with experimental outcomes. Solving for Bayesian Nash equilibria in asymmetric auctions is a notoriously difficult problem that, to our knowledge, had not previously been done in a multiple-unit setting. We developed a new technique for solving close approximations to these equilibria. The basic idea to the approach is to impose a flexible yet constrained parametric offer function for each seller type and then select parameterizations of those functions that minimize the expected gains of a randomly chosen seller unilaterally deviating from their constrained strategy. We were able to find such constrained strategies that brought these expected gains from unilateral deviation extremely close to zero (ie., true equilibrium). Furthermore, the approximate equilibrium strategies predicted actual offers better than reduced-form regressions of offers on observed cost. This approach to solving asymmetric auctions may be useful in many applications.

One potentially interesting application for auctions with quota could be for procurement of environmental or ecosystem services, such as carbon sequestration. Such procurements now take place on a global scale, with a few large-scale buyers buying from many heterogeneous sellers that differ in clearly discernable ways. For example, European brokers may purchase carbon offsets from farmers in Brazil who offer to sequester carbon by returning cropland or pasture to forests. Opportunity costs of providing such services surely vary widely across sellers in ways that European buyers may have a difficult time ascertaining. By placing a quota on the share of offers accepted from any observationally similar group of sellers (say, a given land disposition within a province or county), it will encourage more competition among those sellers, thereby enticing them to make offers closer to their opportunity costs. Without such a quota, sellers would be enticed to make offers near their expectation of the marginal offer accepted, which may be far above their costs.

Although our analysis considers a situation with just two types of sellers, in such an application it is easy to see how the auction might generalize to a situation with many types, and lead to proportionately much greater savings for the buyer. The key feature to notice is that in auctions with quota, offers from a given type are generally less than the highest cost possible from sellers of that type. Thus, if we imagine an environment with a larger number of types–say tens or even hundreds–spanning observable differences in location, then costs

may vary more between types than they do within types. In this kind of situation, significant savings for the buyer might be achieved with minimal loss of efficiency using quotas that limit acceptance to all but a small share of sellers from each type.

A quota mechanism like the one just described would help to address a problem widely described as "non-additivity" in the payments for ecosystem services literature [22, 23, 24]. Non-additivity involves payments for ecosystem services that would have been rendered even in the absence of payment, i.e., at zero cost. Thus, with arms-length buyers and unobservable opportunity costs, attempts to make payments for ecosystem services "non-additive" amounts to an effort to price discriminate across sellers of those services. We have shown that auctions with quota are a simple and effective way of discriminating in this way.

With the idea of looking toward larger scale auctions, such as the global purchase of carbon offsets, future work may further explore whether quota might increase the efficiency of procurement. Auctions without quota are not strictly efficient in an asymmetric setting due to the non-monotonicity of offers that comes from the "weakness leads to aggression" phenomenon, wherein buyers of a higher-cost type make offers below those of lower-cost types but having the same privately observed cost—a phenomenon we observe both in theory and in the laboratory. A modest quota that just restores monotonicity of offers might therefore increase efficiency while simultaneously reducing the buyer's total payment for a given amount of good or service purchased. While the quota selection mechanism involves its own inefficiencies, these could be modest in a large-scale auction wherein a highly-effective quota may need only curtail a small percentage of sellers from any single observable type.

- [1] R. Myerson, Optimal Auction Design, Mathematics of operations research 6 (1981) 58.
- [2] O. Armantier, J. Florens, J. Richard, Approximation of Nash equilibria in Bayesian games, Journal of Applied Econometrics 23 (2008) 965–981.
- [3] J. Bulow, J. Roberts, The Simple Economics of Optimal Auctions, The Journal of Political Economy 97 (1989) 1060–1090.
- [4] B. Kirwan, R. Lubowski, M. Roberts, How Cost-Effective are Land Retirement Auctions? Estimating the Difference Between Payments and Willingness to Accept in the Conservation Reserve Program, American Journal of Agricultural Economics 87 (2005) 1239.
- [5] I. Ayres, P. Cramton, Deficit Reduction Through Diversity: How Affirmative Action at the FCC Increased Auction Competition, Stanford Law Review 48 (1996) 761–815.
- [6] T. Denes, Do Small Business Set-Asides Increase the Cost of Government Contracting?, Public Administration Review 57 (1997).
- [7] P. Milgrom, Putting Auction Theory to Work, Cambridge University Press, 2004.
- [8] R. McAfee, J. McMillan, M. Whinston, Multiproduct Monopoly, Commodity Bundling, and Correlation of Values, The Quarterly Journal of Economics 104 (1989) 371–383.
- [9] E. Krasnokutskaya, K. Seim, Bid Preference Programs and Participation in Highway Procurement Auctions, University of Pennsylvania, Working paper, available at http://bpp. wharton. upenn. edu/kseim/UserFiles/Media/auctsm\_0507. pdf (2008).
- [10] J. Marion, Are Bid Preferences Benign? The Effect of Small Business Subsidies in Highway Procurement Auctions, Journal of Public Economics 91 (2007) 1591–1624.
- [11] T. Hubbard, H. Paarsch, Investigating Bid Preferences at Low-Price, Sealed-Bid Auctions with Endogenous Participation, International Journal of Industrial Organization 27 (2009) 1–14.
- [12] V. Flambard, I. Perrigne, Asymmetry in Procurement Auctions: Evidence from Snow Removal Contracts\*, The Economic Journal 116 (2006) 1014–1036.
- [13] A. Corns, A. Schotter, Can Affirmative Action be Cost Effective? An Experimental Examination of Price-Preference Auctions, American Economic Review 89 (1999) 291– 305.

- [14] R. J. Weber, Multiple-Object Auctions, in: M. S. Richard Engelbrecht-Wiggans, R. M. Stark (Eds.), Auctions, Bidding and Contracting: Uses and Theory, New York University Press, 1983.
- [15] A. Ortega-Reichert, Models for Competitive Bidding under Uncertainty, Ph.D. thesis, Stanford University, Department of Operations Research, 1968.
- [16] W. Gayle, J. Richard, Numerical Solutions of Asymmetric, First-Price, Independent Private Values Auctions, Computational Economics 32 (2008) 245–278.
- [17] M. Rothkopf, A Model of Rational Competitive Bidding, Management Science 15 (1969) 362–373.
- [18] J. Kagel, J. Richard, Super-Experienced Bidders in First-Price Common-Value Auctions: Rules of Thumb, Nash Equilibrium Bidding, and the Winner's Curse, Review of Economics and Statistics 83 (2001) 408–419.
- [19] O. Armantier, J. Richard, Computation of Constrained Equilibrium in Game Theoretic Models: Numerical Aspects, Computational Economics 15 (1997) 3–24.
- [20] V. Krishna, Aunction Thoery, Academic Press, 2002.
- [21] J. H. Kagel, Auctions: A survey of experimental research, in: J. H. Kagel, A. E. Roth (Eds.), The Handbook of Experimental Economics, Princeton University Press, 1995, pp. 501–585.
- [22] P. Ferraro, R. Simpson, The Cost-Effectiveness of Conservation Payments, Land Economics 78 (2002) 339.
- [23] P. Ferraro, Asymmetric Information and Contract Design for Payments for Environmental Services, Ecological Economics 65 (2008) 810–821.
- [24] S. Wunder, The Efficiency of Payments for Environmental Services in Tropical Conservation, Conservation Biology 21 (2007) 48–58.

#### Appendix A. Experimental Auction Instructions and Design

Here we screen shots from our custom-made experimental computer software. Figures A.9 through A.12 show the on screen instructions. These describe how the auction works and how subjects submit offers. Instructions are also read out loud and participants have an opportunity to ask questions. After going through the instructions, there is a demo round, followed by a debriefing screen that quizzes the subject to be sure they understand the rules. A subsequent screen corrects their mistakes and explains why they may have answered incorrectly.

Figure A.13 shows offer screens for each kind of auction in our experiments (without and with quota, respectively). Note that the subject can see the outcome from previous round and can click to open a screen that summarizes outcomes from all earlier rounds.

Figure A.9: Instruction Screens 1–3

Welcome!
Today you will be participating in an experiment on economic decision making
Just for participating in this experiment, you will receive a 45.00 participation payment. However, you can earn substantially more by actively taking part in the experiment. Your total carnings will be paid to you in each at the conclusion of the experiment.
Today we will be conducting about 50 auctions, one after the other.
<ul> <li>Each auction will take about 1 minute.</li> <li>Each auction is separate and independent. Your <i>bids</i> and earnings in any one auction will have no influence on your earnings in any other auction</li> <li>In each auction, you and all other participants are sellers and there is a single buyer. The single buyer is a computer.</li> <li>The buyer will be purchasing takets, which you will be given, in each auction.</li> </ul>
Auction Instructions: Summary
During an auction, each one of you will have one or more tickets that you may offer for sale to the buyer.
You can <i>offer</i> each of these tickets to the seller by submitting a <i>bid</i> .
<ul> <li>The lower your bid the higher the chance your bid will be accepted.</li> <li>If your bid is accepted, you will receive your bid minus the ticket's cost</li> </ul>
Each of your tickets will have a cost . These costs will be randomly determined separately and independently for each ticket in each auction in the experiment.
Both the costs and your bids will be denominated in a currency called E-Bucks. Each E-Buck will be converted into dollars at the end of the experiment at a rate of 0.05 real dollars per E-Buck.

#### **Types of Tickets**

There are several types of tickets, which we call Type A , and Type B tickets.

- For Type A tickets, cost may be any amount between \$0.00 and \$100.00,
  for Type B tickets, cost may be any amount between \$50.00 and \$150.00,

Each of you will have 1 tickets.

Some of you may have a *Type A* ticket
Others of you might have a *Type B* ticket For example ...

No one will have two tickets of the same type.

Before each auction you will be given new tickets, with a new randomly chosen costs, regardless of whether or not you sold any of your tickets in the previous auction.

In other words, each auction is truly independent. Your actions in any auction will have no bearing on your earnings in any other auction.

### Figure A.10: Instruction Screens 4-6

	More on Ticket Types
Even though you wil For example,	l receive new tickets before each auction, your ticket <i>Types</i> will remain the same. If you had a <i>Type A</i> you will always be get a <i>Type A</i> ticket If you had a <i>Type B</i> you will always be get a <i>Type B</i> ticket
The random <i>cost</i> for beginning of each tin	each of your tickets, in any round, will be known to you and only you. Do not let anyone else see them. You will learn the random costs for each of your tickets at the ned round.
Although the <i>costs</i> o	feach of your tickets differs from other participants in the room, the buyer values all tickets equally. To the buyer, any one ticket is as good as any other ticket.
Review	
<ul> <li>There are 10</li> <li>Each particip:</li> <li>There are 2 T <ul> <li>5 of yc</li> <li>5 of yc</li> </ul> </li> <li>Each ticket ha The cost, rar</li> </ul>	total participants (you plus 9 others). ant will recieve 1 tickets. 'ypes of tickets: u have a <i>Type A</i> ticket. u have a <i>Type B</i> ticket. as a <i>cost</i> printed on it. Idomly drawn from an interval that depends on the type of ticket, will be different for every auction.
	Bids and Ticket Scores
During each auction ;	you will have a chance to submit any bid you choose by clicking in the Your bid textbox on the screen and using your number keys.
For each ticket that y	rou submit, a score will be computed. This score is used to rank the tickets with the lower the score the higher the rank.
Once all ticket offers	are received, their scores will be used to determine who the winners are.
The buyer	will accept the 6 lowest scoring tickets scores in each auction.
Tou could have zero	, one of your tickets accepted.
	Maximum Bid
Each ticket is assigr	ned a <i>Maximum Bid</i> . You will <b>not</b> be allowed to submit a bid higher than the ticket's <i>Maximum Bid</i> !
Note that for each t	icket
<ul> <li>Sometimes th</li> <li>Sometimes th</li> <li>Sometimes th</li> </ul>	ne <i>Maximum Bid</i> is <b>much larger</b> than the <i>cost</i> (for example, \$55.00 vs. \$12.00). ne <i>Maximum Bid</i> is a somewhat <b>larger</b> than the <i>cost</i> (for example, \$21.00 vs \$18.00). ne <i>Maximum Bid</i> is <b>less</b> than the <i>cost</i> (for example, \$18.00 vs \$24.00).
Review	

- To offer a ticket you enter a bid
  Each ticket has a maximum bid -- your bid can not exceed this maximum
  A score is computed for each ticket. The buyer will accept the 6 lowest scoring tickets in each auction.

#### Figure A.11: Instruction Screens 7–9



The practice auction, and all other auctions are timed.	<ul> <li>When the timer runs out, any ticket with a non-empty value in its Your bid textbox will be submitted for you automatically.</li> <li>You may submit your offers at any time before the timed round expires, if you wish.</li> <li>There is a timer on the offer screen to assist you in budgeting time.</li> </ul>
---	---

You may not converse with other participants during the entire experiment.

Figure A.12: Instruction Screens 10 and Post-Demo Debriefing

That's it!

In a few minutes, your computer will display a practice screen.

Until then, if you wish you can review the entire introduction.

Please note that a variety of on-line help will be available throughout this experiment -- you can review these rules whenever you want.

#### Earnings: none yet

#### The Demo is complete

Before showing you whether your tickets were accepted, please answer the following questions. After you submit your answers, you will be told the correct answers.

These questions are meant to increase your understanding of how the experiment works. Whether you answer them correctly, or incorrectly, will have <u>no</u> effect on how much you will get paid.

If you do <b>not</b> submit a <i>bid</i> , your earnings for the auction will be zero.	True C	False	еO
If you submit a <i>bid</i> , and it is rejected, the amount you earn for the auction is \$5.00	True 🤇	Fals	eО
Your ticket's cost will be the same in every auction.	True 🤇	False	eО
If you submit a <i>bid</i> and it is accepted, your earnings for the auction will equal your <i>bid</i> minus your ticket's <i>cost</i> .	True C	False	e O
If your <i>bid</i> is less than your ticket's cost, and your bid is accepted, the difference (converted from E-bucks to dollars) will be subtracted from your \$5.00 participation bonus.	True C	False	e O

Submit your answers

### Figure A.13: Seller's Offer Screens

## Without Quota

Results for all the prior auctions	Summary of this auction's rules		Earnings: \$67.60	
Auction in Progress: 6 Your name: A			Time remaining 02:40	Submit your tickets
	Current	auction		Prior auction 5
Type A Ticket Cost \$26.60	45. <u>Your Bid</u>	45 <u>Score</u>	\$18.40 Earnings compute	\$17.00 67 67.00 <u>Eam Score maxScore</u>
_	Au Tabut Tara Davas Canad			Prior auction earnings: \$17.00 Details
There are 10 participants in this auction A total of 6 tickets will be accepted. <u>View details on ticket types</u>	A 0 to 100 B 50 to 150	View prior auction results (all	tickets)	

# With Quota

Results for all the prior auctions	Summary of this auction's rules	<u>Earnings:</u> \$86.00	
Auction in Progress: 7 Your name: A		Time remaining 02:45	Submit your tickets
	Current auction		Prior auction 6
Cost:     \$87.40       TypeLimit:     4	Your Bid Score	Earnings compute	\$18.40 45 45.00 <u>Eam Score maxScore</u>
			Prior auction earnings: \$18.40 Details
There are 10 participants in this auctio A total of 6 tickets will be accepted. However a maximum number of each be accepted. <u>View details on ticket ty</u>	on. <i>Type</i> of ticket can <b>B B C D C C C C C C C C C C</b>	View prior auction	ı results (all tickets)

#### Appendix B. Algorithm for Finding MEGD Offer Strategies

The heart of the procedure for finding approximate equilibria involves calculating the expected gains from unilaterally deviating from an a given constrained parametric strategy. In principle, the constrained strategy may be of any form. For the main part of our analysis we use restricted cubic splines with six knots plus an intercept for each type of seller. We also forced the offers from Type B sellers to equal cost at the highest possible cost draw, which is implied by theory. Because restricted cubic splines require two less parameters than the number of knots, this functional form involved five parameters for Type A sellers plus four parameters for Type B sellers. Given any particular selection of parameters, the expected gains from unilateral deviation are calculated as follows:

- For each type of seller, develop an index across the range of possible cost outcomes. Our index included 100 vales that were equally spaced from the lowest to highest possible outcome. Note that for non-uniform cost distributions, distribution quantiles would be preferred to equally-spaced values.
- 2. Simulate 10,000 auctions by pseudo-randomly drawing costs for each type of seller and apply the constrained strategy as parameterized.
- 3. For each seller type, estimate the probability that any offer will be accepted using the pseudo-random draws for N-1 sellers. Note that because each seller knows their own cost draw, the probability density function of acceptance is a function of the N-1 other sellers, not all N sellers. Furthermore, note that this density is somewhat different for each type of seller since each type possesses different information. We estimate this density using the R package *logConDens* developed by Kaspar Rufibach and Lutz Duembgen. This package computes the maximum likelihood estimator of a non-parametric density that is assumed to be log-concave.
- 4. Given the estimated density functions, calculate the optimal offer for each index value of cost for each seller. Using the optimal and constrained offer for each index value, calculate the gains from unilateral deviation from the constrained strategy for each cost index value for each type of seller.
- 5. Calculate the average of deviation gains across all types and cost index values. This gives the expected gains from unilateral deviation.

The remainder of the algorithm mainly involves minimization of the function described above. The steps are as follows:

1. Pick a functional form for the offer function for each type.

- 2. Select an initial guess for the parameters in the offer function by assuming optimal strategies from a similar symmetric auction.
- 3. Use a Nelder-Mead algorithm to find the parameters that minimize the expected gains from unilateral deviation.
- 4. We re-optimized several times, using the solution from the previous optimization as starting values for the next, repeating until further objective minimization declined by less than 2% of the objective value.