Auction Design with Avoidable Fixed Costs: An Experimental Approach

Wedad Elmaghraby[1]

The Robert H. Smith School of Business, The University of Maryland,

College Park, MD 20742 welmaghr@rhsmith.umd.edu

Nathan Larson

Department of Economics, The University of Virginia,

Charlottesville, VA larson@virgina.edu

Abstract

Advances in information technology and computational power have opened the doors for auctioneers to explore a range of auction formats by considering varying degrees of *bid expressivity* and different *payment rule*, e.g., single price vs. discriminatory prices. While it is clear that one *can* design more complicated auctions, it is still not clear if *should* do so and *which* auction parameters have the greatest impact on the performance on cost and efficiency. The purpose of this paper is to gain some insight into this question, via analytical and experimental methods.

Keywords: Auctions, Experimental, Procurement, Synergies, Asymmetric Bidders

1 Introduction

In recent years, Fortune 500 companies have saved millions of dollars through the increasing use of procurement auctions (Elmaghraby (2006)). At the same time, academic research has produced sophisticated new auction designs tailored to handle complications that are common in procurement. In choosing among these auction designs, a firm may be making implicit choices about how to handle complexity in its procurement problem. For example, some auctions may make the mechanics of the auction more onerous for suppliers (by requiring highly detailed bids) or for the buyer (by requiring a difficult computational problem to be solved in order to choose the winners). Other designs may force bidders to solve more challenging strategic problems in order to bid well. In this paper, we focus on one particular complication – suppliers who have avoidable fixed costs of production – that is commonplace in a procurement context but is often assumed away by

[1]This paper has benefited greatly from the comments provided by seminar participants at UCLA, University of Connecticut, Penn State, Columbia and Stanford University. Both authors gratefully acknowledge support from the National Science Foundation provided under grant DMI 02-05489. The first author also acknowledges support under NSF grant ECCS 0618225

traditional auction models. Using laboratory experiments, we compare the performance of three representative auction formats, each of which handles the complexity introduced by avoidable fixed costs in a different way. We are interested in exploring the following questions:

• How well or how poorly do simple bidding, pricing and allocation rules align individual profit maximizing behavior with the buyer's goal(s) of maximizing productive efficiency and/or minimizing total cost?

• Given the varying complexity of the bid decision problem across auction formats, do subjects behave as equilibrium theory predicts? (That is, do they maximize expected profits with respect to the bidding strategies used by their opponents?) And if not, can we explain why subjects exploit profit opportunities more effectively under some auction formats than others?

To address these questions, we introduce a model in which several suppliers compete to provide a single buyer with multiple units of a homogeneous good (Section 4). Two key features of this model are that the suppliers have private information about their costs, and some suppliers have decreasing average costs, while others have increasing average costs. In the context of this model, an auction design determines how the suppliers compete. To be precise, an auction design specifies the format of the bids submitted by suppliers and the procedure by which those bids are used to determine the quantity that each bidder must supply and the amount that it will be paid. We introduce three auction formats that are intended to capture a design trade-off between 'simplicity' and 'expressivity' (Section 4.2), and use lab experiments to compare the performance of these formats.

Our analysis of the results takes a two-pronged approach. First we identify systematic and significant differences between formats for outcome measures of productive efficiency and buyer cost (Section 5). Then we then characterize approximate equilibrium bidding behavior under our three auction formats [2] in order to compare these equilibrium predictions with the experimental behavior that we observe (Section 6). We find that some of the systematic differences in auction performance are explained by theory, but in other cases, observed behavior diverges markedly from the theoretical predictions. Where there is a gap between the predicted behavior and observed behavior, we assess possible explanations, including alternative models of bidder utility (such as risk or loss aversion) as well as features of the auction formats that make it difficult for bidders to optimize successfully (section 7).

To preview some of the leading results, we find that,

[2] The details, which involve a combination of derived analytical solutions, computed numerical solutions, and approximations, are discussed later.

• An expressive bidding format with discriminatory prices achieves the highest productive efficiency (Section 5).

• A simpler (less expressive) bidding rule with uniform pricing has no pure strategy equilibrium, achieves the lowest productive efficiency of the three formats, and yet also achieves the lowest procurement cost for the buyer. The first two results can largely be attributed to the fact that valuable information is lost when a supplier condenses two dimensions of costs into a one dimensional bid. Given this, the low procurement cost is rather surprising, and is only possible due to extremely competitive bidding in this particular uniform price format (Section 5) .

• Deviations from equilibrium bidding cannot be consistently explained by risk or loss aversion. We offer evidence that these deviations are better explained by differences in the difficulty of the strategic problems that bidders must solve under different formats. This suggests that strategic complexity is an important factor in auction performance, and that equilibrium predictions should be treated with considerable caution as auctions grow more complex (Section 6).

• Surprisingly, deviations from equilibrium behavior sometimes improved rather than detracted from efficiency (Section 6). This occurred because subjects tended to overlook profit making opportunities that *required* coordinated changes in bids under expressive bidding (Section 7.2.2).

In the next two sections, we briefly discuss the relevance and implications of avoidable fixed costs, often appealing to wholesale electricity markets as a leading example, and discuss the placement of our work relative to the existing literature. Then we describe our model, auction formats, and experimental procedure, followed by analysis of the experimental results. We conclude with some suggested directions for future work.

2 Avoidable Fixed Costs

One of the key challenges in designing auctions for B2B settings is that suppliers' costs are often comprised of both avoidable fixed as well as variable costs. For example, in the wholesale electricity markets, generators incur start-up costs if they are called upon to generate. Avoidable fixed costs are also present in most manufacturing and service settings (e.g., the start up costs associated with running a production line for a customer's specific order and/or training employees.) The creeping use of auctions into 'new' B2B markets, such as electricity, manufacturing, brings the role of avoidable fixed costs and their inclusion in the auction design discussion front and center.

Unfortunately, many of the standard results of auction theory do not necessarily apply when bidders have avoidable fixed costs. For example, many conclusions about bidding strategy depend on the fact that, under standard pricing rules, a supplier who bids along its marginal cost curve will always cover its total cost.[3] When bidders have avoidable fixed costs, this is no longer true – suppliers can lose money even if their variable costs are covered at the margin. On a related point, when production costs are comprised of only variable costs, the efficient allocation of production can be determined simply by ranking suppliers by marginal cost. However, in the presence of avoidable fixed costs, efficiency requires considering the bidders' entire cost curves. (That is, not only marginal costs, but all inframarginal costs need to be assessed.) To illustrate the practical importance of these complications with a concrete example, we turn to a brief description of wholesale electricity auctions (O'Neill et al. (2005)). Billions of dollars are at stake in these markets, and market participants are acutely (and vocally) aware of the complexity of designing an auction where some bidders have decreasing average costs. [4]

The California Power Exchange was the first auction market for wholesale electricity within the US, and allows generators (the suppliers) to bid in variable (energy) cost only. The auctioneer then 'stacks' the bid curves to create an aggregate (increasing) supply curve, and the market is cleared at the price that balances supply and demand. This uniform price is paid to all suppliers including those with offers below the clearing price. While this auction format has the benefit of simplicity of design, a potential drawback is that bid expressivity is limited – the format does not provide any way for a bidder to directly express its start-up costs (among other limitations). In principle, this means that it may be difficult for a supplier to bid competitively while also shielding itself against the risk of making a loss (if it is dispatched for less than its intended quantity). In the literature, this type of risk is often described as an exposure problem, and the associated losses may be described as confiscatory. In practice, it has been argued that the exposure problem in the CALPX auctions is mitigated by the uniform price format and the fact that the suppliers with the most significant start-up costs (such as nuclear and coal plants) tend to be inframarginal. Since these suppliers will tend to receive a market clearing price well above their variable costs, their prospects for recovering those start-up costs may be fairly good. The single uniform price may also be helpful to bidders in their planning: *"... Indeed, the clearing-price auction is an essential feature of any electricity market designed to reliably provide consumers electricity at minimum cost. The clearing-price auction plays a critical role in the least-cost scheduling and dispatch of resources, and provides an essential price signal both for short-run performance and long-run investment incentives." (Cramton and Stoft (2007)).* The question of how these potential advantages and

[3] One may think of uniform or discriminatory (pay-as-bid) pricing rules, for example.

[4] We fully acknowledge that we abstract away from many of the salient features of an electricity market that further complicate the auction design, for example the presence of must-run generation plants that displace inframarginal bids. However, the discussion below does not suffer from their absence, since our focus is on the importance of bid expressivity and payment rules on bidder behavior.

drawbacks to the CALPX format balance out in practice is far from settled.

In contrast to California, Pennsylvania-New Jersey-Maryland (PJM), which operates the largest competitive wholesale electricity market in the world, has approached its design problem from a traditional 'optimization' perspective. The PJM market allows suppliers to clearly express their complex cost structure by submitting start-up (an avoidable fixed cost) as well as variable costs, i.e., submitting a multi-part bid. PJM then runs a large Mixed Integer Program to determine the generators' dispatch based on this more complete information about costs. From this optimization program, the auctioneer identifies a uniform (single) clearing price with *individualized (discriminatory)* side-payments, used to ensure that none of the dispatched suppliers are producing unprofitably given the uniform price and reported start-up costs. While the multi-part bid structure allows bidders to better express their underlying cost structures, together with the allocation optimization program, it tends to make the relationship between market inputs (bids) and outputs (prices and allocations) somewhat opaque. Some observers worry that this might encourage bidders to devise complicated schemes to try to manipulate the market.[5] Other observers argue that even well-intentioned bidders may find it difficult to understand what the allocation optimization program can tell them about their most appropriate placement in the market. In addition, some auction theorists have advised against using a discriminatory pricing format for electricity, expressing concern that the multitude of price signals may create confusion and lead to more of a 'guessing the market price game' (Kahn et al. (2001)).[6]

These two contrasting approaches to auction design with avoidable fixed costs set the scene for our research plan.

3 Position within Existing Literature

In selecting a procurement auction's format, a buyer would ideally design a mechanism that aligns individual bidders' profit maximization behavior with her own objective(s). To do so, the buyer must understand the impact her selected auction format has on a supplier's submitted bid. Auction theorists have been able to identify relatively simple auction formats (simple bidding formats,

[5]For example, some European markets only allow generators to submit their start-up costs bids on an annual basis - while permitting daily bids for variable cost. These auction designers believe that this asymmetric bid flexibility will help reduce any strategic bidding along the start-up cost dimension (Harbord et al. (2002))

[6]In 2001, California hired a Blue Ribbon panel to explore if switching to a discriminatory auction from a uniform-price auction would help improve performance of CA-PX (Kahn et al (2001)). The panels conclusion was that a move to discriminatory pricing may create more problems than it solves. Under incomplete information, a discriminatory pricing rule creates a Guess the Clearing Price game; could result in inefficient dispatches due to errors in guessing.

pricing and allocation rules) that align individual profit maximizing behavior with the buyer's desire for efficiency and/or minimization of expected total cost (Myerson (1981)).

When suppliers have multiple cost components , a buyer must consider the impact of the *mechanical expressivity* of the bid structure on the submitted bids and on the efficiency of the allocation rule. That is, does the bid structure allow the supplier to fully express his cost structure (if he were to choose to do so) or does it require him to condense his cost information into a less expressive bid format? Clarke (1971) and Groves (1973) demonstrate that if suppliers are able to submit bids over all possible subsets of goods, then there exists a pricing and allocation rule for which each supplier will submit his true costs in an auction; that is, no supplier misrepresents his costs, given the bid structure. Note that the full expressivity required under their proposed mechanism (referred to as a Vickrey-Clarke-Groves mechanism) requires the bid and pricing space to be as large as 2^N when there are N items being auctioned, which may be impractical in many settings (Hobbs et al. (2000) and Rothkopf (2007)). If suppliers are provided with anything less than full expressivity in favor of simpler bid and/or pricing formats, it is not always clear how they will behave.

There has been a long history of testing the predictions of auction theory via controlled experiments with human subjects (please see Kagel (1995) for a overview). Much of this literature addresses the standard setting in which bidders desire at most one unit and performed tests of the Revenue Equivalence Theorem, Winner's Curse and the presence of collusion. Our paper differs from this body of work in that we consider a setting where bidders desire more than one unit and may have superadditive preferences.

Motivated by the recent use of combinatorial auctions in transportation, research has emerged that addresses the design of *iterative* auctions for multiple objects when bidders desire more than one object or unit (e.g. Ledyard et al. (2002) and Kwasnica et al. (2005)). In these auctions, bidders are allowed to submit bids on different packages of goods/units. The ability to condition their bids on the composition of the package allows bidders to, in theory, accurately reflect any superadditive preferences they may have while avoiding the exposure problem (Plott (1997)), which arises when a bidder is not able to obtain the entire package of items desired. To the best of our knowledge, our work is the first to analyze the impact of package versus non-package bid formats in *simultaneous* (sealed-bid) auctions. Furthermore, we consider varying the payment structure and its impact, both theoretically and experimentally, on the performance of the auction.

This work falls within the emerging cross-disciplinary area of *behavioral operations*; papers in this area verify/compare via controlled human experiments the predictions of normative models in operations. Topics addressed have included the newsvendor model (Bolton and Katok (2005) and Schweitzer and Cachon (2000)), the beer game (Croson and Donohue (2002) and Wu and Katok

(2006)) and supply chain coordination (Cui et al. (2007)). Our work adds to this growing area by providing a systematic analysis of observed and equilibrium behavior under various procurement auction formats, and provides arguments that it is the way in which the suppliers' costs interact with institutional and strategic features of an auctions to make it difficult for otherwise rational, risk neutral subjects to be effective profit-maximizers.

4 The Model, Auction Formats, and Experimental Procedure

4.1 Model

Consider a single buyer who demands $D = 3$ discrete units of a good. If necessary, it can produce the good in-house at constant marginal cost $R = \$100$; D and R are commonly known constants.

The buyer faces 3 suppliers. Each supplier can produce either 0, 1, or 2 units of the good. A supplier either has increasing, constant or decreasing average costs. Suppliers with decreasing average costs may be thought of as facing a fixed cost of producing a positive quantity (which can be avoided by producing 0) and a constant marginal cost. We capture this range of supplier types using a family of cost structures given below, where c_q is a supplier of type θ's average cost to produce q units,

$$C(0) = 0 \qquad c_1 = 100 - \theta \qquad c_2 = \frac{100 + \theta}{2}$$

Each supplier independently and privately draws a different cost parameter θ from the distribution $\theta \sim U(0, \bar{\theta})$, where $\bar{\theta} = 50$. Note that the average cost is increasing for $\theta \in [0, \frac{100}{3})$, decreasing for $\theta > \frac{100}{3}$ and constant for $\theta = \frac{100}{3}$. We will often refer to a supplier type by its cost to supply one unit, $c_1 = 100 - \theta$. Notice that c_1 moves inversely to θ : low c_1 suppliers are more efficient at producing one unit and high c_1 suppliers are more efficient at producing two units. The cost structure is carefully designed so that no supplier type dominates any other type (in the sense of having lower costs at every production level).

This setting is designed to capture, as simply as possible, certain features that we would like to study: (i) A mix of efficient production scales. Low type suppliers ($c_1 < 66\frac{2}{3}$) minimize average cost by producing one unit, while high type suppliers ($c_1 > 66\frac{2}{3}$) minimize average cost by producing two units. (ii) Overall production generally involves multi-sourcing and asymmetric production levels for the winning suppliers. This follows from the fact that no single supplier can provide the full three units demanded, so an allocation will typically procure the full capacity of two units from one supplier and the remaining one unit from another supplier.[7]. (iii) An efficient overall production

[7]Production of one unit by all three suppliers is also a possibility, but given the avoidable fixed costs, it is rarely efficient.

schedule typically requires both high and low type suppliers to participate. This follows from the point immediately above, and the fact that high c_1 suppliers are absolutely (not just relatively) the least cost producers for two units, and conversely for low c_1 suppliers and one unit. (*iv*) Efficiency cannot be determined at the margin. Due to the presence of avoidable fixed costs, an efficient production schedule must consider both marginal *and* inframarginal costs. (*v*) We parameterize supplier types with a one-dimensional index, θ. This is a technical point, but an important one – avoiding a multi-dimensional type space allows for a much more complete theoretical analysis than would otherwise be possible and also permits us to make relatively more observations per type in the lab.

4.2 Auction Formats

In our experiments we consider three auction formats, 1U, 2U and 2D, that highlight certain design trade-offs and are close analogues to commonly used formats in industry.[8] 1U is a format in which bidders submit *one*-part bids and there is a *uniform* clearing price. As we shall see, adequately conveying one's willingness to supply with a one-part bid is challenging for a supplier with avoidable fixed costs. Format 2U retains a uniform clearing price, but increases bid expressivity by permitting suppliers to submit *two* part bids. However, given the lumpy nature of production, competition is effectively occurring at two margins: competition to be the larger, two unit supplier, and competition to be the smaller, one unit supplier. A single uniform price cannot hope to make both of these margins transparent. Format 2D retains two-part bidding, but switches to *discriminatory* pricing, giving bidders feedback in the form of separate (and potentially different) prices on each of the three units. Moving from 1U, through 2U, to 2D we progressively increase the flexibility of bidding and the richness of the feedback to bidders. At the same time, the sheer volume of the information required from, and provided to bidders increases as well. Below we describe the auction formats more precisely and present some illustrative examples.

1U : One-part bids, uniform pricing (1U)

Under 1U, each bidder i simultaneously submits a one-dimensional bid b_i. This bid is a pledge to make its entire capacity, or any part of it, available at any price per unit greater than or equal to b. The buyer procures the entire capacity of the lowest bidder (in this case 2 units from the lowest bidder) and the remaining one unit from the next lowest bidder. All bidders are paid a market-clearing/uniform price per unit equal to the bid of the marginal bidder. For example, consider the following three submitted bids from suppliers 1, 2, and 3 respectively, $b_1 = 77,\ b_2 = 85,\ b_3 = 80$.

[8]Deriving optimal sales mechanisms is not our purpose here, and we make no claims about the optimality of any of these formats.

Given these bids, the quantity allocations would be, $Q_1 = 2,\ Q_2 = 0$ and $Q_3 = 1$ and the price paid per unit would be \$80.

2U : Two-part bids, uniform pricing (2U)

Under 2U, the bidding space is expanded to reflect the supplier's cost structure, but all units continue to be paid the same uniform price. Each bidder i submits a two-part bid (b_{i1}, b_{i2}) indicating the minimum price *per unit* it is willing to be paid if it supplies one unit (b_{i1}) or two units (b_{i2}). The market-clearing price is defined to be the lowest price at which it is possible to procure exactly three units by fulfilling some subset of the submitted bids, taking at most one bid from each supplier. For example, if the following three bids were submitted, $(b_{11}, b_{12}) = (60, 77)$, $(b_{21}, b_{22}) = (70, 67)$, and $(b_{31}, b_{32}) = (100, 55)$, the market-clearing price would be 60 per unit. At a price of 60, Supplier 1 is willing to produce one unit and Supplier 3 is willing to produce two units. At any lower price, it would be impossible to procure more than two units. If there is only one subset of bids for which it is possible to procure three units at the market-clearing price – as is the case here – then those bids form the allocation. In this example, b_{11} and b_{32} are accepted, and the suppliers are obligated to produce quantities $Q_1 = 1$, $Q_2 = 0$, and $Q_3 = 2$ respectively. If there is more than one subset of bids that would provide a total of exactly three units at the same market clearing price, then the winning allocation is the one for which the total cost would be lowest if each unit were paid as bid.[9]

Notice that although the spirit of the uniform pricing rule is similar to 1U – that is, all units are paid the lowest price at which it is possible to procure three units – in this format, the determination of the price and the winning bids cannot be accomplished as simply as it was under 1U. The buyer must now solve an (albeit simple) integer program to find the price and assignment of production that satisfy the conditions above. Furthermore, this format may require the buyer to pass over some bids below the market-clearing price (b_{22} in the example above) if they are attached to quantities that do not fit well with the buyer's needs and the other accepted bids.

2D: Two-part bids, discriminatory pricing (2D)

Bidder i submits a two-part bid (b_{i1}, b_{i2}) as under 2U. However, under 2D, each supplier is paid the amount of its own accepted bid per unit supplied. A supplier's bids are still mutually exclusive: i will either have b_{i1} accepted, supply one unit, and be paid b_{i1}, or it will have b_{i2} accepted, supply two units, and be paid a total of $2b_{i2}$, or it will have no bid accepted and produce zero. The rules specify that the buyer will accept the combination of bids that (i) includes at most one bid from

[9]For example, with these submitted bids, b_{11} and b_{32} will be accepted; although b_{11} and b_{22} yield the same cost to the buyer.

each supplier, (ii) totals exactly three units of supply, and (iii) minimizes the total payment that the buyer must make. For the earlier example of $(b_{11}, b_{12}) = (60, 77)$, $(b_{21}, b_{22}) = (70, 67)$, and $(b_{31}, b_{32}) = (100, 55)$, the allocation remains $Q_1 = 1,\ Q_2 = 0,\ Q_3 = 2$, but now Supplier 1 is paid 60 per unit, while Supplier 3 is paid 55 for each of two units. At the conclusion of the auction, all of the bidders see the price vector (55,55,60).

4.3 Experimental Design

In the Fall of 2005, we ran a series of experiments in the Netcentricity Behavior Lab at the University of Maryland (UMD). Each session focused on one of the three auction formats and included 12-16 subjects, all of whom were undergraduates at UMD. Subjects were informed that they would be paid a 'show-up' fee of \$10, plus any profits that they earned during the session. Earnings during the session were measured in an experimental currency ('francs') and were later converted to dollars at the rate of 50 francs to \$1 USD. The conversion rate was selected so that, based on equilibrium predictions, subjects would earn roughly \$20[10].

In our experiments, we discretize the type space described earlier. A supplier could be one of seven types, evenly spaced between $\theta = 0$ and $\theta = 50$. Given the inevitability of a certain amount of background noise in experimental behavior, we believe that seven types suffice as a reasonable approximation to opponent types drawn uniformly from $[0, 50]$. At the same time, capping the number of types at seven allows us to obtain an adequate number of observations (observed bidding behavior) per type. Figure 1 lists the types, their average costs for producing 1 and 2 units, respectively, as well as plots the average cost curves.[11]

Upon arrival, subjects were given written instructions. After reading these, they were led through a Powerpoint presentation illustrating the auction procedure and the computer screens they would encounter. We then ran three practice rounds and stopped to answer any remaining questions before beginning the experimental auction rounds.

The experimental phase of each session consisted of 30 rounds. In each round subjects were randomly matched in groups of three to compete in an auction. Each subject was informed of his own costs, but not those of his opponents. Subjects kept the same costs for six rounds at a time,

[10] Losses during the session were deducted from the show-up fee. While some subjects did make overall losses during the auction rounds, these were never larger than the show-up fee, so there were no "bankruptcies."

[11] The types are only approximately spaced between $\theta = 0$ and $\theta = 50$ because costs are rounded to integers. Note also that there are no "round" numbers in the cost table (multiples of 5 or 10) by design. Our pilots indicated that round costs tended to produce bids that were also anchored to multiples of 5 or 10, an artifact that we wanted to avoid.

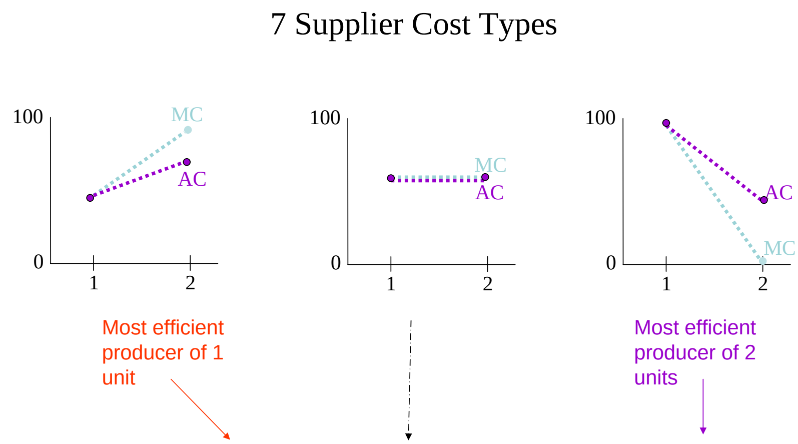

c_1	50	58	67	75	83	92	100
$c_2 = C(2)/2$	75	71	67	63	58	54	50

Figure 1: Seven possible supplier cost types in experiment. The top graphs plot the average costs for three supplier types, at output levels of 1 and 2.

drawing new costs in Rounds 1, 7, 13, 19, and 25.[12]. We will refer to each of these blocks of six rounds as a *sub-session*. At the end of each round, subjects were shown: (*i*) Their costs, (*ii*) Their bid(s), (*iii*) How many units they won in that round, (*iv*) the winning price(s), (*v*) their profit in that round, and (*vi*) their cumulative profit including the $10 show-up fee. Each session ran for about 90 minutes, start to finish, and subjects were paid immediately upon completion of the session.

5 High Level Results

In this section we present summary results from the three treatments: 1U, 2U, and 2D (Table 1). We will begin with a high-level comparison of total supplier cost (SC) and buyer cost (BC) across the auction formats. SC, which is the total cost to suppliers of providing the three required units, provides a measure of the efficiency of the auction format. Meanwhile, BC measures the total

[12] In order to ensure that every cost type was represented, a different type was assigned to each of the first (randomly ordered) seven subjects, then to subjects 7 through 14, and so on. Subjects in excess of a multiple of seven were assigned costs randomly. In practice, this means that a subject's chance of facing an opponent with the same type was lower than if types had been fully random. This should be taken under consideration (along with other approximations that will be made) when the equilibrium predictions are presented. However, this has no bearing on the analysis of computed best responses, since these are based on the actual distribution of opponent bids.

payments made by the buyer. In a private sector procurement context, minimizing buyer cost may be the principal issue for auction design, while in markets with regulatory oversight, efficiency (minimizing SC) may be an equal or greater concern.

Table 1: Summary Statistics on Buyer's Cost and Suppliers' Cost under 1U, 2U and 2D.

	Treatment			
	Random	1U	2U	2D
% efficient	24%	33%	68%	77%
$\frac{(SC - SC_{eff})}{SC_{eff}}$	16.7%	10.1%	2.9%	1.9%
SC		186.7	173.7	170.1
BC		187.5	212.2	188.1
$\frac{(BC - SC)}{SC}$		0.4%	22.2%	10.6%

For each instance of an auction in the lab, we define an efficient production schedule to be an allocation of supply to the three bidders that procures three units at the least total (true) cost. Then, SC_{eff} is defined to be the average total production cost of an efficient schedule across all of the auctions in a treatment. Thus, SC_{eff} is a lower bound on the average supplier cost that could potentially be achieved in a treatment. While the *expected* value of SC_{eff} is the same across all of the treatments, the realized values will be slightly different due to random variation in the matching of cost types. In each case, "Random" refers to an allocation in which one supplier is randomly selected to supply two units and another to supply one unit. While perfect efficiency provides a best-case benchmark for each auction format, one can think of the random outcome as a worst-case benchmark; any auction design that performs close to, or worse than, a random allocation is doing poorly indeed. We look at two measures of productive efficiency: the frequency of efficient production, and the average gap between actual and efficient supplier cost. The first measure simply computes the fraction of auctions in a treatment for which the actual supplier cost is equal to the efficient (i.e., lowest achievable) supplier cost. The second measure is defined as $\frac{SC - SC_{eff}}{SC_{eff}}$; that is, the "inefficiency markup" of the actual average supplier cost of the suppliers selected via the auction, SC, over the efficient average cost SC_{eff}.

For each of these two measures, there is a clear and consistent ranking: two-part bidding with discriminatory prices performs the best, while two-part bidding with uniform pricing does somewhat less well. In each case, one-part bidding with uniform pricing performs relatively poorly. Note that $\frac{(SC - SC_{eff})}{SC_{eff}}$ is smallest for 2D, and largest for 1U.

Next we turn to a comparison of buyer cost across the formats. Buyer cost may be treated as the sum of two terms: supplier cost, and a "markup" term. Thus buyer cost is affected by two factors: the efficiency of production (SC), and the competitiveness of the auction format ($BC - SC$). The buyer's procurement cost is lowest with one-part bidding and uniform pricing (1U), closely followed by two-part bidding and discriminatory pricing (2D); two-part bidding with uniform pricing (2U) is substantially more costly. It is striking that the least efficient (by far) auction format, 1U, is also the cheapest procurement option for the buyer. This is only possible because the average supplier profit margin under 1U is almost nonexistent, as the markup column of the table indicates.

To assess the significance of these differences in average efficiency and buyer cost across the three auction formats, we ran several simple pooled regressions with treatment variable dummies. In these regressions, a unit of observation is a single auction; we have a total of 420 observations (150 under 1U and 2D, and 120 under 2U). The left-hand side variables are *Efficient* (equal to 1 if the auction allocation is efficient and 0 otherwise), *Markup* (equal to $BC - SC$, loosely, the suppliers' profit margin), and *Buyercost* (equal to BC). In the basic specifications, the right-hand side variables include the cost types of the three suppliers, where *c1min*, *c1med*, and *c2min* are defined to be the lowest one unit cost, median one unit cost, and lowest two unit cost, respectively. Note that whenever a two supplier production schedule is efficient (and this is almost always the case), the efficient total cost is $SC_{eff} = \mathit{c1min} + 2\,\mathit{c2min}$, which does not depend on the median supplier at all. However, we are completely agnostic about how these cost types might affect realized efficiency and profit margins. Also included are *period* (the auction round, between 1 and 30) and *modperiod* (the round of the sub-session, from 1 to 6), to capture learning and any other time trends. Finally, we have indicator variables $1U$, $2U$, and $2D$ (with $1U$ generally omitted). These regressions were run using OLS, with the exception of *Efficient*, for which a probit was used. (Further details about these specifications, as well as alternative regression specifications, are reported in online appendices B and C.)

The basic specifications are reported in Column 1 of the tables in online appendix C. The results corroborate our summary statistics. In the *Efficiency* probit, the coefficients on $2U$ and $2D$ are both positive and significant at the 1% level, indicating that two-part bidding is associated with a greater likelihood of efficient allocations. Furthermore, the additional advantage of $2D$ over $2U$ is significant at the 10% level. However, as the *Markup* regression shows, higher efficiency is balanced by significantly (and substantially) higher supplier profit margins under the two-part bidding formats, particularly for $2U$ (the difference between $2U$ and $2D$ is significant at the 1% level). The combination of these two effects is evident in the *Buyercost* regression: there is no significant difference in the buyer's total procurement cost between $1U$ and $2D$, while $2U$ is significantly more expensive. For $2D$, the benefit of improved efficiency (i.e., reduction in $SC - SC_{eff}$) is roughly absorbed by the suppliers (as an increase in $BC - SC$), leaving the buyer no better off. For $2U$, the reduction in competitiveness vastly exceeds the efficiency gain, leaving the buyer substantially worse off.

In order to more clearly trace out the relationship between auction formats, our subjects' bidding behavior, and key outcome measures like efficiency and buyer cost, it helps to break the rest of the analysis up into two parts. In the first part (Section 6), we characterize the equilibrium behavior in these markets, that is, we take it as given that subjects are rational, risk-neutral profit maximizers. Embodied in rationality are the ideas that subjects perfectly understand the strategic environment, develop correct beliefs about the bidding of opponents, and do not consistently miss opportunities to improve their payoffs. This allows us to make a "best-case" evaluation of how successfully each auction provides incentives that align individual profit-seeking with efficient dispatch and competitive bidding.

As we will show in table 2, our subjects' bidding departs substantially from equilibrium predictions. The second part of our analysis (Section 7) attempts to understand why subjects may be more successful in exploiting profit opportunities in some settings and less successful in others.

6 Equilibrium Analysis

Before we characterize the approximate Bayesian Nash equilibria of the three auction formats to be studied in the lab, we should take a moment to explain why these equilibrium predictions are approximate and argue that these approximations are entirely adequate given the limited role that the equilibrium analysis is intended to serve.

One approximation, which carries through all of the formats, is that we will analyze equilibria of auctions in which cost types are drawn uniformly from a continuous interval ($\theta \sim U[0, 50]$), while in our experimental treatments costs are actually drawn from a discrete grid of seven values (also on $[0, 50]$). In principle, these are entirely distinct strategic situations, and inferences about behavior in one setting should not be applied to the other. In practice, they are probably closer than they seem. In price-setting models like these with discrete types, and in the absence of noise, pure strategy equilibria tend to fail because "cliffs" in the probability distribution of rival bids create incentives to very slightly undercut those rivals – leaving only mixed strategy equilibria. However, in an experimental setting, there is always a certain amount of white noise in actions, implying that the probability distribution of rival bids is rather smooth, as it would also be if types were continuous. Simulations buttress this argument – for each of the auction formats, noisy best response dynamics with discrete types tend to converge close to the equilibria that we present for continuous types. The other approximations that we make have to do with the computation of equilibria: in one case, our analytical solution is approximate (2D), while in another (1U) numerical computation of the equilibrium is required. Here too, simulations suggest that the error involved is minor.

Table 2: Equilibrium Predictions versus Observed Behavior under 1U, 2U and 2D.

	Treatment		
	1U	2U	2D
% efficient	33%	68%	77%
% efficient$_{eqm}$	37%	41%	97%
BC	187.5	212.2	188.1
BC_{eqm}	193	289	200
$\frac{(BC-SC)}{SC}$	0.4%	22.2%	10.6%
$\frac{(BC_{eqm}-SC_{eqm})}{SC_{eqm}}$	5%	58%	17%

Table 2 summarizes the equilibrium predictions about efficiency in the three treatments and the actual levels of efficiency achieved by the subjects. Actual efficiency levels are neither as good as theory would predict for 2D (and to a less extent 1U), nor as poor as theory would suggest for 2U. The first half of this observation has a familiar flavor – we are accustomed to seeing auction formats where observed behavior falls short of equilibrium predictions (Kagel (1995)). However, the second part of the observation is more puzzling – under auction format 2U, deviations from equilibrium behavior tend to improve rather than detract from efficiency. Below we summarize the derived equilibria for each treatment; formal details are in online appendix A. We then describe and offer explanations for these deviations in section 7.

6.1 1U

Under 1U, each bidder places a single bid. The bids are stacked from lowest to highest, and supply is allocated first to the lowest bidder (up to its limit of two units), then to the next lowest bidder, who supplies the marginal third unit. Bidders are paid a uniform price equal to the marginal (second lowest) bid for each unit procured.

Before considering equilibrium bids, notice that the 1U allocation scheme cannot avoid a degree of *institutional* inefficiency: even if we were free to dictate the bidding strategies used by suppliers, we could not achieve the efficient dispatch all of the time. The problem is that any set of bidding strategies assigns a fixed order according to which different cost types will be dispatched under the allocation rule. But no single fixed dispatch order is efficient for every realization of cost types. Recall that c_1 is inversely proportion to a supplier's average cost for two units, or equivalently, its relative advantage in producing at a larger scale. By construction, no supplier has an absolute cost advantage; those that are relatively more efficient at producing two units (high c_1 types) are relatively less efficient at producing one unit, and *vice versa*. Given this, the efficient allocation will always require the highest c_1 supplier to produce two units and the lowest c_1 supplier to produce one unit, while the supplier in the middle (who is not the least cost producer at any production scale) should be out of the market. Realized efficiency in the market will depend on how consistently the ranking of bids matches this c_1 cost ordering. For example, on one day, the bidders who show up might have cost types (c_1) equal to 100, 92, and 50, while on another day, the types might be 92, 83, and 50. Achieving an efficient allocation on both days would require the bidding strategies to satisfy both $b(100) < b(50) < b(92)$ (for the first day) and $b(92) < b(50) < b(83)$ (for the second day), which is impossible. Here, an allocation rule based on ranking one-dimensional bids is simply not nuanced enough to track allocations that depend sensitively both on the full set of bidders and on both dimensions of each bidder's costs, which is what efficiency would require.

As a measure of sheer institutional inefficiency, we calculated the expected production cost under the best fixed dispatch order relative to the expected production cost under an efficient allocation. We found that the best expected production cost that can be achieved with one-dimensional bids is about 8% higher than the cost of the efficient allocation. The bid ordering (according to types) that achieves this second-best level of efficiency is $d^* = \{100, 92, 83, 58, 50, 75, 67\}$ - it requires low bids from the highest c_1 types, moderate bids from the lowest c_1 types, and high bids from 'middle' c_1 types.

Equilibrium Bids

Let us turn to equilibrium bids. Any pure strategy equilibrium would imply a fixed dispatch order over the cost types. One natural (but inefficient) order would be the one that ranks supplier types from highest to lowest (in terms of c_1). Another, (less natural but more efficient) order might inject the suppliers who are most efficient at marginal production into the middle of the order, in the manner of d^*. Of course, there are many other possibilities. A necessary condition for any of these orders to represent an equilibrium is that each supplier's position in that order is incentive compatible. We would typically settle for checking that incentive compatibility is satisfied locally – that is, no cost type can improve its expected payoff by bidding *slightly* higher or lower in order to move up or down in the dispatch order. With avoidable costs, expected payoffs may not be single-peaked, so we also need to check incentives *globally*. For example, an intermediate cost type may have two potential bids – bid low and aim to be inframarginal or bid high and aim to poach one unit at the fringe of the market – each of which is locally optimal, but only one of which is a global best response.

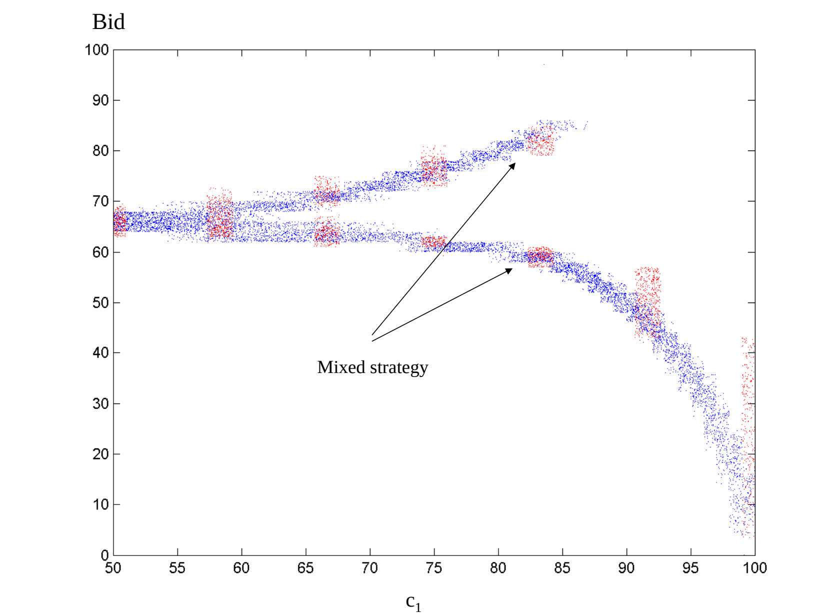

Figure 2: Mixed equilibrium bidding strategy under 1U (blue (red) dots indicate simulation with continuous (seven discrete) types.

Our analysis indicates that these local and global incentives interact in such a way to preclude the existence of *any* pure strategy equilibrium – any equilibrium involves mixing by some cost types. We have approximated the equilibrium using numerical simulation methods; the results, presented below, provide considerable intuition about the incentives at work in this auction. Figure 2 summarizes the equilibrium strategies. The scatter plot shows the frequency of different bids for each cost type after many rounds of simulation.[13] As the figure indicates, the highest type suppliers ($c_1 > 85$) and the lowest type suppliers ($c_1 < 60$) appear to have pure strategies involving low and moderate bids respectively. So far, this is consistent with the constrained-best dispatch order suggested by d^*. However, intermediate cost types sometimes bid high and sometimes bid low. To understand why this mixing is inescapable in equilibrium, consider what would happen if these intermediate types always took the upper or the lower "branch" in Figure 2. Bidding according to the upper branch is akin to poaching at the fringe of the market – these suppliers will be out of the dispatch most of the time and will occasionally sell one unit and set the price. But if *all* intermediate cost types were

[13] Bids from the transient initial phase of the simulation are not included, so the plot represents the distribution of bids after strategies have settled down.

to bid this way, there would be less competition inframarginally and the market price would tend to be higher, making it too tempting to submit a low bid and gamble on being the inframarginal supplier. However, to have any reasonable chance of being inframarginal, most of these types would have to bid below c_1, their cost of supplying a single unit. As long as few other intermediate types are bidding low, the chance of having the second lowest bid and having to produce one unit at a loss is acceptably low. However, if all of these intermediate types were to bid low, then many more of them would wind up marginal (and making losses), while prices would be lower for those lucky enough to be inframarginal. On average these types would then lose money and would do better to make modest profits with a high fringe bid. In equilibrium, there must be just enough competitive bids from these suppliers that they are indifferent between bidding competitively or at the fringe.

6.2 2U

In this format, each bidder places two bids, $b_1(c_1)$ and $b_2(c_1)$, where $b_i(c_1)$ indicates the lowest per unit price at which supplier of type c_1 it is willing to supply exactly i units. The uniform market-clearing price p is set equal to the lowest price at which it is *possible* to procure exactly three units by accepting at most one bid from each supplier. By soliciting two bids from each supplier, the auctioneer could in principle (i.e., setting aside bidder incentives) elicit enough information to always allocate supply efficiently, avoiding the institutional inefficiency of format 1U. And for our cost structure, format 2U does in fact allocate production efficiently *if bidders were to submit bids equal to their actual costs.*[14] However, strategic incentives drive bidding in a very different direction. Below we present equilibrium bidding strategies under 2U for the version of our model with continuous costs. Simulations suggest that this equilibrium is both unique and a qualitatively accurate approximation of equilibrium with discrete cost types.[15]

The symmetric equilibrium bids, as a function of one-unit cost c_1 are

$$\begin{aligned} b_1(c_1) &= \begin{cases} 3c_1 + 100\ln(1 - c_1/100) + 100\ln 3 - 100 & \text{if } c_1 \leq \frac{200}{3} \approx 66.7 \\ 100 & \text{if } c_1 > \frac{200}{3} \end{cases} \\ b_2(c_1) &= 0 \end{aligned}$$

The proof that these strategies constitute an equilibrium is in online appendix A; we provide below intuition behind the equilibrium.

[14]This is not a general feature of 2U – there are other cost structures for which it cannot allocate efficiently all of the time.

[15]Specifically, we wrote Matlab code to simulate a version of best response learning with discrete cost types. The dynamics consistently converged close to the equilibrium presented here. Since the equilibrium is not one that subjects would be likely to adopt by introspection alone, the fact that it appears to be globally stable under learning is reassuring.

Note that $b_1(c_1)$ is greater than 90 for the $c_1 = 50$ and is increasing in c_1, so bids for one unit are quite high. Meanwhile, second unit bids are pooled at zero. To better understand this 'highball-lowball' bidding strategy result, observe that if these strategies are followed, the price will always be set by a (very high) one unit bid, and the lowest two unit bidder will supply two units at this price. Given these strategies, the price will be high enough that it is attractive to all cost types – even $c_1 = 50$ – to try to submit the lowest b_2 and sell two units. Since submitting a lowball b_2 bid is essentially costless – in equilibrium it never turns out to set the price – this undercutting incentive sends b_2 down to the lowest permissible bid, which is zero. Why then doesn't competition do more to bring down the extremely high one unit bids? Since a one unit bid end up being price-setting, a supplier faces the standard first price tradeoff: by increasing its bid slightly, it increases its profit in the event that it sells one unit, at the cost of being out-of-market more often. However, there is an additional factor here: an increase in b_1 that takes a supplier out of the running to supply one unit may put it back in the mix to supply two units. Conversely, bidding b_1 more competitively risks cannibalizing the chance to sell two units. Since two unit profits are so fat, the disincentive to bid b_1 competitively is quite strong, thereby reinforcing the reason that two unit profits were so large in the first place. For a large fraction of the low type (c_1) suppliers, this incentive induces them to pool together at 100 – the least competitive bid that the buyer allows.

Because of the complete pooling of two unit bids and the partial pooling of one unit bids, the equilibrium allocation is predicted to be efficient only about 41% of the time, and the average supplier cost in equilibrium is approximately 6.7% above its efficient level. This represents only a slight improvement over the predicted efficiency under format 1U. Given the discussion above, it should be no surprise that the competitiveness of pricing under 2U is quite poor: the expected markup of total cost to the buyer over total cost to the suppliers is 58%.

6.3 2D

As in format 2U, here each supplier places two bids, $b_1(c_1)$ and $b_2(c_1)$. However the payment and allocation rules are different. A supplier of type c_1 dispatched for i units is paid its bid of $b_i(c_1)$ per unit, and the winning bids are those that minimize the total cost to the buyer while supplying exactly three units. As with 2U, allowing bidders to place a one unit bid and a two unit bid makes it possible, in principle, to elicit all of the cost information that would be needed to achieve an efficient allocation. However, in contrast with 2U, bidders' strategic incentives under 2D tend to promote efficiency rather than undermining it.

The bidding strategies below are an exact equilibrium for the model with continuous cost types, under the assumption that only two bidder allocations are permitted. The strategies are only approximately optimal if three bidder allocations are permitted, but because three bidder allocations only occur in around one in thirty auctions, the approximation is quite close. Numerical simulations confirm that these strategies are a good representation of equilibrium in the auction with discrete cost types and a small amount of noise.

$$\begin{aligned} b_1(c_1) &= \frac{100}{3} + \frac{2}{3}c_1 \\ b_2(c_1) &= \frac{11}{12}100 - \frac{1}{3}c_1 \\ &= 25 + \frac{2}{3}c_2 \end{aligned}$$

These bidding strategies leverage a feature of our model that is convenient but not particularly general. Under our cost structure, the ranking of types by one unit and by two unit costs are mirror images of each other: the lowest cost one unit supplier is the highest cost two unit supplier, and so on. This turns out to imply that, in equilibrium, a supplier is never marginal for both one unit dispatch and for two unit dispatch at the same time. That is, a supplier's adjustments to b_1 only matter in realizations for which this bid is in, or is close to being in, the market dispatch. But in precisely these events, the supplier's two unit bid b_2 is *not* in (or close to being in) the market. The reverse of this statement – adjustments to b_1 are only relevant in events in which b_2 is not close to being competitive – is also true. As a result, a supplier can get away with approaching its two bids as separate and unrelated profit maximization problems. It is almost as if there were two separate markets – one to procure a block of two units and the other to procure one unit. This separation features prominently in the equilibrium strategies: one can confirm that each strategy $b_i()$ above represents the symmetric equilibrium of a stand-alone three bidder auction in which a block of i units is to be procured.

The cost structure also implies that, if a two supplier dispatch is efficient, then this efficient allocation should involve the lowest (highest) type supplying one (two) units. Because $b_i(c_1)$ increases monotonically in c_1, the equilibrium allocates supply efficiently in all of these cases. Where the equilibrium comes up short is in some realizations with three low type suppliers. In these events it would be efficient to procure one unit from each bidder, but because these bidders mark up b_1 more than b_2, an allocation with one of them producing two units turns out to be cheapest for the buyer. These cases are relatively uncommon: the equilibrium allocation is efficient approximately 96.5% of the time, and the average supplier cost is only 0.2% above its efficient level. Competitiveness of this format lies in between the other two auctions considered, with an equilibrium expected total cost to the buyer that is 17% above the expected total production cost to the sellers.

7 Disequilibrium Analysis : Best Response Behavior

As we saw in table 2, there is a substantial divergence between equilibrium predictions about the efficiency of each format and experimental outcomes. In the remainder of the paper, we describe these deviations, argue that they can be attributed to imperfections in profit maximization, and assess explanations that could explain these imperfections. Loosely, one might categorize these explanations according to whether they relax the "profit" or the "maximization" side of the standard model of behavior. That is, one class of explanations would assert that subjects are successful maximizers, but that their utility functions include objectives other than profits (reducing risk or losses, for example). Another class of explanation would be that subjects do care about profits, but face challenges in trying to maximize those profits. We find little support for the first class of explanation and considerable support for the second. We will argue that the *ways* in which subjects stumble when confronted with challenging maximization problems are both systematic and instructive.

Comparing a subject's actual bidding strategy to an equilibrium bidding strategy is not a particularly useful way to evaluate whether the subject is profit-maximizing effectively, since the equilibrium strategy is only a best response to *equilibrium* bidding by opponents (but not necessarily to their *actual* bidding). A better comparison, which we pursue below, requires computing best responses (BR) to the actual strategies used by opponents. The basic approach is: for each auction format, and for each cost type, compute the expected profit that this type would earn from every possible bid *if* its opponents' bids were drawn randomly from the full sample and with the associated frequency of bids faced by that cost type, for that particular treatment[16].

Table 3: Suppliers responses to profit making opportunities.

	1U	2U	2D
Actual Avg. π	0.26	13.02	5.99
Average potential π	3.45	25.53	9.07
% π left on table	92%	51%	33%
Shortfall	-3.3	-12.5	-3.1

Table 3 summarizes the average actual supplier profits, relative to the average best response profits, for each auction type.[17] Subjects leave money on the table in every case, and they do particularly

[16] Interested readers can contact the authors for the code and data used to generate these Best Response Bids.

[17] These are the weighted averages of the average profit for each cost type and the best response profit for each cost type, respectively. The weights correspond to the frequency with which each cost type appeared in the treatment.

poorly under 2U, where the profit opportunities are largest. In the rest of this section, we explore potential explanations for why profit opportunities are not fully exploited. These can loosely be grouped into explanations based on *alternative preferences*, i.e., subjects may not be risk neutral profit maximizers, and explanations based on *learning*, i.e., strategic or institutional factors may make profit opportunities relatively more difficult to identify under certain formats.

7.1 Alternative Risk Preferences

7.1.1 Risk Aversion

Risk aversion is often invoked in the context of single unit first price auctions to improve the fit between theoretical and empirical bidding behavior. Loosely, a risk averse subject facing an uncertain outcome between either losing the auction (earning 0) and winning (earning a positive surplus) will be willing to trade off more of that surplus to improve its chance of winning than a risk neutral subject would. Consequently, equilibrium bidding by risk averse subjects is more competitive than it would be if subjects were risk neutral. Explanations favoring risk aversion generally suffer from the critique that the supporting evidence, usually overly competitive bidding, is also consistent with many alternative explanations. One advantage of our data is that we are able to study settings in which risk aversion and overly competitive bids are not synonymous.

Consider the case of the 1U format. Notice that 1U has a bit of the single unit first price flavor, since by raising its bid, a marginal, price-setting bidder increases the payment it receives if it remains marginal but also increases the chance that it fails to supply any units at all. However, bidders more generally face *quantity risk*, as any given bid may result in a bidder being obligated to supply zero, one, or two units. Bidding more competitively always increases one's chance of being inframarginal (and hence being required to supply two units) and generally has a mixed effect on the chance of supplying the single marginal unit. (Starting from a high bid, reductions tend to move bidders who are out-of-market onto the margin, but below a certain level, further reductions tend to move bidders at the margin inward, reducing their chances of supplying one unit.) The result is a complex balance of risks; for example, for a cost type that is most efficient at producing one unit, it is not immediately clear whether this balance favors bidding higher (risking being out of the market) or bidding lower (risking winning two units, which it might have to produce at a loss). Furthermore, bidders who would like to be inframarginal face a combination of quantity risk and *price risk*. Bidding high enough to cover one's two-unit cost may carry an unacceptable risk of accidentally winning only one unit (which would need to be produced at a loss), while bidding

(Given the process by which types were generated, this is close to (but not identical to) equal weighting.)

below two-unit cost improves the chance of winning the desired two units but introduces a chance of receiving a price below that two-unit cost – again making a loss. This rich set of tradeoffs allows us to test the implications of risk aversion in a much more nuanced way than is typically possible.

Our angle of attack is to revisit our computation of the expected profit function for each cost type, relative to the empirical distribution of bids. However, now, rather than compute the expected profit at each potential bid, we compute the bidder's expected utility under the assumption of either constant absolute, or constant relative risk aversion. Some illustrative results are presented in Figure 3. (Further details are available upon request.) For each cost type, the columns represent the risk neutral best response bid, the median bid among our subjects, and the best response bid for a subject with CRRA preferences.[18] One should bear in mind that the results have not been smoothed, and therefore retain some of the idiosyncracies of the sample data. Nonetheless, the pattern is striking. In two cases, the risk averse and risk neutral best responses are identical, or nearly so, but in the other five cases, the risk averse best response is higher, so risk aversion implies bidding that is *less* competitive than the risk neutral level. However, in five out of seven cases, actual bids are *more* competitive than the risk neutral level. Turning to the exceptions, for $c_1 = 50$, risk neutral, risk averse, and actual bids are all very close, providing no positive evidence for risk aversion. And for $c_1 = 100$, while actual bids are modestly less competitive than the risk neutral best response, this is a drop in the bucket; a bidder who is even slightly risk averse would vastly prefer to be out of the market entirely with a bid of 100.

While Figure 3 presents these results for a particular CRRA specification, they are extremely robust to alternative parameters and CARA specifications. Intuitively, payoff variability is so substantial for bids that are deep in the market – both because of quantity and price risk – that any risk averse supplier, even a large one, does best by nibbling at the margin with a relatively high bid. For example, for cost type 75, the standard deviation of profits is about three times higher (12 times expected profits, vs. four times expected profits) when it chooses the lower bid rather than the higher bid in its mixed strategy. And for cost type 83, the standard deviation of profits is about 17 times larger at the lower of the bids over which it mixes![19] This completely fails to describe the behavior that we see among subjects in 1U. Since 1U is the treatment that arguably exposes subjects to the most risk, risk aversion, at least as it is usually modeled, does not provide

[18]For this computation, the coefficient of relative risk aversion was set to 0.5, roughly in line with parameter estimates obtained in the experimental auction literature. Initial wealth was assumed to be fixed at 500, the initial balance in subjects accounts at the start of the treatment. The broad pattern of results that we cite is not at all sensitive to changes in these parameters.

[19]Comparing these standard deviations to expected profits is not particularly meaningful for type 83 because those expected profits are essentially zero.

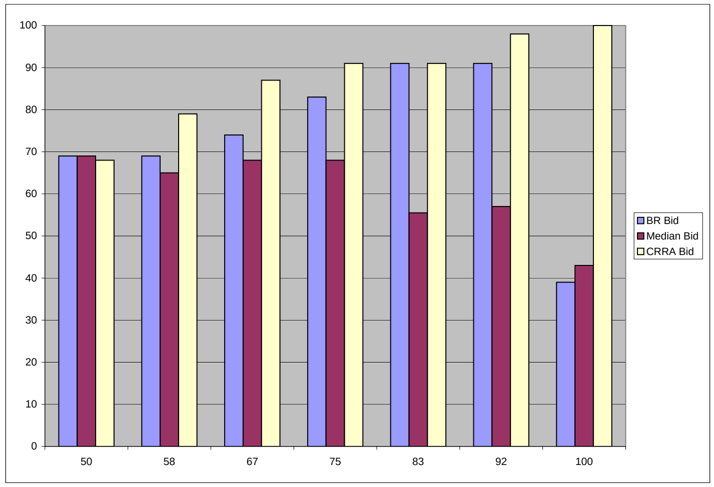

Figure 3: Best Response bids under risk-neutrality (BR Bids), constant relative risk aversion (CRRA bids) compared with the median bids under 1U.

a consistent explanation for deviations from profit maximization among our subjects. That does not mean that we should give up on risk and losses entirely as explanatory factors, but any role that they play will need to be relatively nuanced.

7.1.2 Loss Aversion

There is now substantial evidence that people do not treat gains and losses symmetrically (Kahneman et al. (1991)). This is manifested in a number of types of behavior which collectively have been termed "loss aversion." We take a relatively reduced form approach to studying loss aversion among our subjects. In two of our auction formats, 1U and 2U, profit maximization sometimes requires placing bids that could result in losses for some realizations of the opponents' bids.[20] If subjects are not pure profit maximizers but also place some weight on avoiding losses *per se*, then we would expect their actual strategies to expose them to a smaller chance of losses than the best response strategies would have done. For each treatment, and for each cost type, we can identify a set of *loss exposed* bids – these are the bids that could result in a loss for some realization of opponent bids. For 1U, any bid $b < \max(c_1, c_2)$ is in this set, since a subject bidding b could earn

[20]For treatment 2D, this is never the case. Since a bidder can bid for one and two units separately, and a winning bidder is always paid her bid, there is never any incentive to place a bid below cost.

as little as $b - c_1$ (if dispatched for one unit) or $2(b - c_2)$ (if dispatched for two units and the marginal bid is also b). With the two-part bidding of 2U, a bid (b_1, b_2) is loss exposed if either $b_1 < c_1$ (chance of selling one unit at a price less than c_1) or $b_2 < c_2$ (chance of selling two units at a price below c_2). Table 4 presents the fraction of all bids that were loss exposed for each of the two treatments; for comparison we also present the fraction of bids that would have been loss exposed if subjects had used the best response strategies computed earlier.

Table 4: Suppliers' responses to profit making opportunities.

	1U	2U
% Loss Exposed under BR	16%	60%
% Loss Exposed (actual)	77%	12.5%

In isolation, bidding under format 2U could be consistent with a story in which subjects bid so as to avoid the possibility of losses, even if this reduces expected profits. However, bidding in 1U does not fit this story at all – while the best response strategies generally involve fairly high, unexposed bids, subjects actually tend to choose lower bids that both earn lower expected profits *and* expose them to the chance of losses.

Given the evident willingness of subjects in 1U to embrace the possibility of losses, it becomes much more difficult to explain bidding in 2U in terms of loss aversion alone. The difference could be reconciled if the two subject pools differed substantially and systematically in their attitudes toward losses, but this seems unlikely. [21] A basic point that emerges is that it is hard to believe that subjects with a clear understanding of their prospective payoffs could bid so aggressively/competitively under 1U and yet so timid under 2U. If one is willing to assume that the subject pools are similar in their attitudes toward losses, then the bidding in 2U becomes rather puzzling. As a result, we are led to consider whether subjects do, in fact, have a clear understanding of their prospective payoffs, and if they do not, what features of the auction formats tend to make profit opportunities harder to find. It is to this issue that we turn next.

[21] To take one illustrative case, consider the type with costs $(c_1, c_2) = (75, 63)$. This type's median bids were $b_1 = 80$ and $b_2 = 70.5$ respectively. It observed an average price of 73.9 in the auctions in which it was present, dispatched one unit (two units) 38.6% (8.3%) of the time, and earned an average payoff of 2.8. As it turns out, the profit-maximizing strategy for this type would have been to highball its one unit bid ($b_1 = 100$), lowball its second unit bid ($b_2 = 1$), and *always* sell two units at whatever price turned out to clear the market. This strategy does incur losses if the market-clearing price happens to be below 63. However, these losses would have been rare and small: a bidder would have made an average loss of about 8.3 about $\frac{1}{7}$ of the time. The gains, on the other hand, would have been substantial: $\frac{6}{7}$ of the time, the bidder would have gained about 33.4 on average. The expected payoff with the profit-maximizing strategy would have been about 27, or close to ten times the actual payoff.

7.2 Explanations Based on Learning and Bidding Heuristics

We might expect that as subjects acquire more time and experience bidding with a particular set of costs, their skill in identifying profit opportunities will improve, and less money will be left on the table. Since the average profits discussed above cover the entire treatment, they might tend to present an excessively pessimistic picture of how successfully subjects are ultimately able to profit-maximize. This is of particular concern, since we have argued that these auction formats are strategically complex and difficult to master. Recall that subjects kept the same cost type for six rounds before randomly drawing new costs. To determine whether subjects were more successful after adjusting to their new costs, average and best response profits were recomputed separately for the rounds at the beginning and at the end of a six round cycle; the results are in Table 4.[22] In the following, we will use the shorthand of referring to profits in Rounds 1-3 and Rounds 4-6 when we really mean all of the rounds that are among the first three or the final three of a six round cost cycle.

The evidence for improvement is mixed. Under auction format 2D, the average profit left on the table falls by about 47% between the first three and last three rounds with the same costs. However, for format 1U, there is only a modest decline, and for 2U, subjects actually forego *more* profit, not less, in the last three rounds. This suggests that deviations from profit maximization cannot be easily dismissed as a transient phenomenon.

Table 5: Suppliers responses to profit making opportunities.

Treatment	1U3		2U		2D	
	1-3	4-6	1-3	4-6	1-3	4-6
π_{Actual}	0.2	0.5	10.8	12.3	5.1	6.7
π_{BR}	4.3	3.9	22.4	25.5	9.6	9.1
Shortfall: $\pi_{Actual} - \pi_{BR}$	-4.0	-3.4	-11.6	-13.2	-4.5	-2.4
% captured	5%	13%	48%	48%	53%	74%

Table 5 focuses on average performance across cost types and does not provide any indication of whether some suppliers adapt better than others. Figure 4 breaks these results down to show the actual and best response profits, at each supplier cost level, for Rounds 1-3 and Rounds 4-6. Several features are worth noting. First, experienced subjects (Rounds 4-6) perform consistently well across cost types in treatment 2D, but this is not true for the other two treatments. And

[22] That is, we compute one set of results for Rounds 1-3, 7-9, 13-15, 19-21, and 25-27, and another set of results for Rounds 4-6, 10-12, 16-18, 22-24, and 28-30.

under 2U, the only substantial improvement comes from the extreme type $c_1 = 100$, while the low c_1 types maintain a stable profit (with relatively little room for improvement), and the performance of interior cost types generally deteriorates, despite substantial unexploited profits. For suppliers whose appropriate production level is unclear, profit exploitation is much more hit or miss and additional experience does not necessarily seem to help. In the rest of this section, we will argue that it is the way that the cost structure interact with institutional and strategic features of auctions 1U and 2U that makes it difficult for otherwise rational, risk neutral subjects to be effective profit-maximizers.

7.2.1 'Marginal Intuition' Bidding Heuristic

In the absence of avoidable fixed costs, observations at the margin provide interested parties (suppliers, the buyer, a social planner concerned with efficiency) with all the information they might need to make optimal decisions. For example, a buyer concerned with efficiency need only verify that marginal cost is equalized across suppliers. Meanwhile, a supplier concerned with its own profit need only be able to determine whether raising its bid a little bit tends to raise or reduce its profit on average. By making small adjustments in the direction of higher profits, it will eventually arrive at its optimal strategy.

With avoidable fixed costs, marginal signals no longer suffice for making good decisions. Assessing the efficiency of an allocation then requires evaluating suppliers' entire cost curves in order to minimize the cost of inframarginal units (a step that happens "for free" when costs are purely variable). For a supplier, this means that experimenting with small changes to its strategy and adjusting in the direction of higher profits is not guaranteed to eventually lead it to an optimal strategy.

How does this issue arise in the context of our specific auction formats? It may help to illustrate the difficulties with 1U and then discuss why 2D seems to perform better. With 1U, one-dimensional bids essentially allow bidders to turn a knob controlling a one-dimensional tradeoff between receiving high prices (bid high) and selling larger quantities (bid low). But suppliers with economies of scale may not have single-peaked preferences over this tradeoff: producing two units at a low price or demanding an out-of-market price and shutting down may both be better than producing one unit at an intermediate price.

This issue is well illustrated by the potential best response profits for two of the large supplier types: (c_1, c_2) equal to $(83, 58)$ and $(92, 54)$, see Figure 5. In both cases, the profit functions have multiple peaks: bidding either quite low or quite high is preferable to intermediate bids. And in both cases, bidders who are new to these costs (Rounds 1-3) tend to get stuck near the low-bidding peak, even though they would have been better off bidding high. What is perhaps more interesting is that as market conditions change (Rounds 4-6), the bidders do not adjust. The type 92 bidders happen to be in the right place at the right time – their low-bid peak turns into the global optimum, and their profits rise with the tide. However, the type 83 bidders are not so lucky – they go down with the ship as profitability at their local peak deteriorates. If anything, their distribution of bids shifts up marginally – this is a shift in the right direction, but something substantially more than a marginal change is needed.

While the evidence is clearly circumstantial, these bidding patterns are consistent with a story in which a supplier adopts new costs with an initial idea about where it expects to fit into the market (either deep in the market or poaching at the margin) and experiments in the neighborhood of that initial guess, but does not manage to identify and exploit profits outside of that neighborhood.

In contrast, format 2D gives suppliers two knobs to turn, controlling two tradeoffs: one for the likelihood of small scale dispatch versus price, and another for the likelihood of large scale dispatch versus a different price. To the extent that these tradeoffs can be regarded as independent, a supplier has single-peaked preferences over each one, and therefore is able to "tune" its one-unit profits and its two-unit profits by making separate marginal adjustments to each bid. In general, this assumption of independence is not warranted – each part of a mar ket allocation may depend on all of the bids, not just a particular subset of them. However, there are special features of our cost structure that make it a reasonable approximation to treat the 2D auction as composed of two separate and independent sub-markets. This issue of the interdependence or modularity of the market allocation is what we will turn to next.

7.2.2 'Modular Intuition' Bidding Heuristic

In a sense, format 1U forces suppliers to try to handle multiple tradeoffs with a one-dimensional bid. On the other hand, formats 2U and 2D provide suppliers with enough instruments (two bids) to handle all of the tradeoffs they face, but in order to bid successfully, they need to understand how those instruments *jointly* influence market outcomes. In certain cases, as we will see with 2D, it may be reasonable to take a modular approach to multi-dimensional bidding by focusing on one market segment at a time and ignoring interdependencies. However, in general, a modular approach to bidding need not be very successful. To illustrate this point, we turn to format 2U.

Recall that the equilibrium for this format calls for highball bids (at or near 100) for one unit and lowball bids (equal to 0) for two units, for *all* supplier cost types. Mechanically, this equilibrium is sustained by the fact that the price is always set at a very high level by some supplier's one-unit bid, making it desirable for even the smallest scale suppliers to bid low for two units with the hope of being inframarginal. Conversely, the incentive to bid more competitively for one unit is undercut by the fact that this would cannibalize the chance of being dispatched for those large two-unit profits – this keeps the one-unit bids high.

It is important to note that this equilibrium is robust – it is the consistent outcome when a dynamic best response process is simulated. However, to reap the full benefits of this best response strategy a bidder must consider changes in both bids at the same time. To see this, suppose a hypothetical disequilibrium situation in which one-unit bids are uniformly spread over $[80, 100]$ and two-unit bids are equal to 10, and consider the strategy of a supplier with cost type $(67, 67)$. Imagine that this supplier has been bidding $(b_1, b_2) = (80, 67)$ and is considering whether it could do better. At the moment, it is always dispatched for one unit (setting the price at 80) and never for two units, earning a profit of 13. Holding b_2 fixed and considering adjustments to b_1 in isolation, it is not difficult to show that this supplier cannot improve on $b_1 = 80$. It will never be considered for two units as long as both other bidders are bidding $b_2 = 10$. Bidding its own b_1 below 80 simply reduces its payment without changing its allocation, and bidding b_1 above 80 (for these parameters) is not worth the chance of being out-of-market.

Alternatively, hold b_1 fixed and consider adjustments to b_2 in isolation. Given the low two-unit and high one-unit bids by opponents, our supplier's one-unit bid of 80 will always be needed in the equilibrium dispatch, regardless of how it adjusts b_2. Thus, isolated changes to b_2 have no effect on this supplier's payoff – it continues to earn 13.

Finally, consider adjusting both bids simultaneously to $(100, 0)$. The supplier's one-unit bid is no longer competitive; the market price is now set by the lesser of the two opponents' one unit bids – in expectation, this is approximately 87. Furthermore, the auction rules give our supplier priority over the other opponent for the inframarginal two units by virtue of our lower two-unit bid. The expected profit earned by bidding $(100, 0)$ is therefore roughly $2\,(87 - 67) = 40$. In this example, there are substantial profit improvements to be made, but they require making a coordinated increase in b_1 and a decrease in b_2; a supplier who makes adjustments one bid at a time would see no evidence that its current strategy is sub-optimal. The specific way that bids interact here in determining the market allocation is obviously particular to this model, but the general idea – that a supplier's own bids may compete with or complement each other in ways that are not immediately intuitive – applies quite broadly.

Figure 6 presents level curves of the 2U profit functions computed with respect to empirical bids for two of the cost types, 75 and 83, that perform relatively poorly. In each case, a black circle marks the location of the median bids used by that type in the treatment.[23] For both cost types, the actual bids are near a "corner" of the level curves: profits are not very sensitive to decreases in b_2 or increases in b_1, but increase dramatically with a shift to the southeast (simultaneously decreasing b_2 and increasing b_1). In other words, the situations are fairly similar to the example discussed above. As in the previous section, the evidence is suggestive rather than conclusive, but it would appear that our subjects are relatively successful in exhausting the improvements that can be achieved by adjusting one component of their strategies at a time, but tend to overlook opportunities that would require coordinated shifts in several components at once. Given the full bid expressivity provided under 2U, this could be viewed as a type of *satisficing* behavior (Radner (1975), Tversky and Kahneman (1981)), whereby subjects bid enough to cover their costs for each quantity, and do not aggressively explore further profit making opportunities. A by-product of this modular behavior is that bidders submit bids that are roughly proportional to their costs, resulting in higher efficiency under 2U than equilibrium predicts.

Why then do subjects fare relatively well in format 2D? As we discussed earlier in section 6.3, our cost structure has the feature that the cost ranking of types for production of one unit is the mirror image of the cost ranking for production of two units. A bidder's optimal choice of, say, b_1 depends on tradeoffs between events in which it has the strictly lowest cost for one unit and events in which it is just tied for the lowest one-unit cost. (Raising b_1 improves its winning price in the first set of events but converts a win to a loss in the second set.) Similarly, its optimal choice of b_2 depends on tradeoffs involving events in which it has the (weakly) lowest two unit cost. But because of the structure of costs, these events never overlap. Thus, adjustments to b_1 are only payoff relevant in events where they cannot affect the bidder's (losing) standing in the two unit auction, and *vice versa*. As a result, bidders are approximately justified in treating their two bids as strategically independent. [24]

8 Conclusion

One of the key challenges in designing auctions for B2B settings is that suppliers' costs are often comprised of both avoidable fixed as well as variable costs. Advances in information technology and computational power make it increasingly feasible to use more expressive bids and more sophisticated allocation and payment rules. However, as auction designs grow more complex, the assumption that bidders are perfect profit maximizers becomes considerably less innocuous. In this

[23] That is, the circle marks $(median\,(b_1)\,, median\,(b_2))$.

[24] The 'approximate' justification arises because (i) costs are drawn from a discrete, not continuous space, and (ii) we neglect the (very small) chance of outcomes in which all three bidders supply one unit each.

paper, we introduce an environment with avoidable fixed costs and auction formats that handle those costs in different ways. Across these auction formats, we study how effectively bidders respond to the individual incentives that they face, and how closely those incentives line up with typical buyer goals of maximizing productive efficiency or minimizing total procurement cost.

Through equilibrium and experimental analysis we find that the 'simplicity' of a one-part bid structure is deceiving: From amongst our three auction formats, 1U posed the most difficult bidding problem for bidders with very small profit-making opportunities for most supplier types. In many practical applications, such as electric power procurement, simple auctions have been criticized for failing to adequately account for complexity in suppliers' costs. Nonetheless, these simple mechanisms often persist. Our results suggest that in certain cases, this persistence can be entirely rational – simple one-part bidding schemes like 1U, despite being terribly inefficient, can nevertheless prove to be very cost competitive for the buyer. Broadly speaking, suppliers with decreasing average costs face the danger of being rationed to one unit at an unprofitable price as the marginal bidder. They can try to avoid this outcome either by bidding more competitively (to be more likely to be inframarginal) or by bidding less competitively (to be sure to cover one unit cost). Our subjects predominantly take the former approach, and as a result, the market-clearing price grows very competitive as the market composition shifts toward high type suppliers.

In contrast, the more complicated (or viewed differently, more expressive) 2D format aligned the buyer's objectives of efficiency and cost with the supplier's individual profit maximizing behavior both theoretically and in the experiments. We conjecture that the structure of 2D allowed subjects to engage in simple bidding heuristics, namely marginal and modular bidding, and still exploit profit making opportunities. Interestingly, 2U performance in the lab was significantly different from its theoretical prediction. The 'highball - lowball' nature of its equilibrium bidding strategy proved difficult for subjects to discover and exploit. Hence, failure of marginal intuition, but ability to employ modular bidding strategies gave rise to a bidding strategies that resemble a satisficing rule of covering costs.

We present this paper as a first step to understanding the impact of greater bid flexibility and payment rules on the resulting bidding behavior. As a first step, we leave many unanswered questions for future research. For example, with 'stable' market participants participating in a repeat auction settings, does the increased communication space offered by multi-part bids create greater concerns for signalling and collusion? We consider a setting where a supplier must select between two possible allocations, one or two units. Furthermore, when ranked from most to least efficient, suppliers had a complete opposite ranking for one unit versus two units. As we consider a more general setting where there are many more possible award scenarios, and hence possible overlap in rankings from one allocation scenario to the next (and hence the separateness of the markets will fail), it is not clear if bidders will be able process the richer market feedback under discriminatory pricing and continue to take advantage of multi-part bids to respond to profit making opportunities. These are just a couple of the directions we hope that future research will take the field as we continue to explore auction design in richer market settings.

9 Bibliography

Ausubel, L. M. and P. Cramton, 2002, "Demand Reduction and Inefficiency in Multi-Unit Auctions", Working Paper University of Maryland

Brenner, D. and J. Morgan, (1997)," The Vickrey-Clarke-Groves versus the Simultaneous Ascending Auction: An Experimental Approach", Working Paper, Princeton University.

Bolton, G.E., Katok, E. (2004), "Learning-by-doing in the newsvendor problem: a laboratory investigation of the role of experience and feedback", Working Paper, Penn State University.

Clarke, E. (1971)," Multi-part pricing of public goods," *Public Choice*, vol. 2, pp. 19-33.

Cramton, P. (2003), "Electricity Market Design: The Good, the Bad, and the Ugly" Proceedings of the Hawaii International Conference on System Sciences, January

Cramton, P. and S. Stoft (2007), "Uniform-Price Auctions in Electricity Markets", *The Electricity Journal*, Vol. 20, No. 1, pp. 26-37.

Croson, R. and K. Donohue (2006), "Behavioral causes of the bullwhip effect and the observed value of inventory information" *Management Science*, Vol. 52, No. 3, pp. 323-336

Cui, T. H., J. S. Raju and Z. J. Zhangy "Fairness and Channel Coordination" *Management Science* forthcoming.

Elmaghraby, W. (2006), "Auctions within the e-Sourcing Process", to appear in *Productions and Operations Management*

Groves, T. (1973), "Incentives in Teams," *Econometrica* Vol. 41, pp. 617-631.

Harbord, D. and N. Fabra and N. von der Fehr, (2002). "Modeling Electricity Auctions," Working Paper 0206001, Game Theory and Information

Hobbs, B. F., M. H. Rothkopf, L. C. Hyde, and R. P. O'Neill (2000), "Evaluation of a Truthful Revelation Auction for Energy Markets with Nonconcave Benefits," *Journal of Regulatory Economics*, Vol. 18, No. 1, pp. 5-32.

Kagel, J. (1995), “Auctions: A survey of experimental research,” in Kagel and Roth eds. *Handbook of Experimental Economics* Princeton University Press, Princeton.

Kahn, A., P. Cramton, R. Porter and R. Tabors, (2001) “Pricing in the California Power Exchange Electricity Market: Should California Switch from Uniform Pricing to Pay-as-Bid Pricing?” Blue Ribbon Panel Report, California Power Exchange, January

Kahneman, D., J. L. Knetsch, and R. H. Thaler (1991), “Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias”, *The Journal of Economic Perspectives*, Vol. 5, No. 1, pp. 193-206

Kwasnica, A.M. Ledyard, J., Porter, D. and DeMartini, C. (2005), “A New and Improved Design for Multiobject Iterative Auctions”, *Management Science*, Vol. 51, No. 3, pp. 419-434

Ledyard, J. O., M. Olson, D. Porter, J. A. Swanson, and D. P. Torma (2002), “The First Use of a Combined-Value Auction for Transportation Services ” *Interfaces*, Vol. 32, No. 5, pp. 4-12

Myerson, R. (1981), “Optimal auction design,” *Mathematics of Operations Research*, Vol. 6, pp. 58-73.

O’Neill, R. P., P. M. Sotkiewicz, B. F. Hobbs, M. H. Rothkopf, and W R. Stewart, Jr. (2005), “Efficient Market-Clearing Prices in Markets with Nonconvexities,” *European Journal of Operations Research*, Vol. 164, pp. 269-285.

Plott, C. (1997), “Laboratory Experimental Testbeds: Application to the PCS Auction” *Journal of Economics and Management Strategy*, Vol. 6, No. 3, pp. 605–638

Radner, R. (1975) “ Satisficing”, *Journal of Mathematical Economics*, Vol. 2, pp. 253-262.

Rothkopf, M. H. (2007) “Thirteen Reasons the Vickrey-Clarke-Groves Process is Not Practical,” *Operations Research*, Vol. 55 , pp. 191-197.

Schweitzer, M.E. and G.P. Cachon (2000). “Decision Bias in the Newsvendor Problem with a Known Demand Distribution: Experimental Evidence”, *Management Science* Vol. 46, No. 3., pp. 404-420

Tversky, A. and D. Kahneman (1981), “The framing of decisions and the psychology of choice,” *Science* Vol. 211, No. 4481, pp. 453-458 .

Wu, Y. and E. Katok (2006), “Learning, Communication, and the Bullwhip Effect,” *Journal of Operations Management*, forthcoming.

10 Appendix: Figures

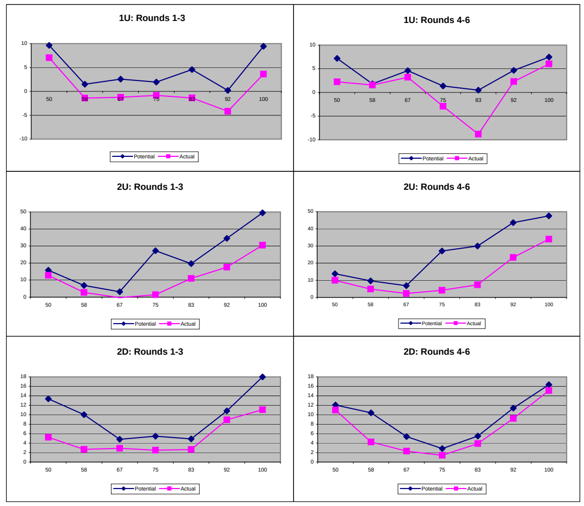

Figure 4: Average supplier profits, by type over the first 3 rounds a subject was assigned a cost, versus the last three rounds for 1U, 2U and 2D.

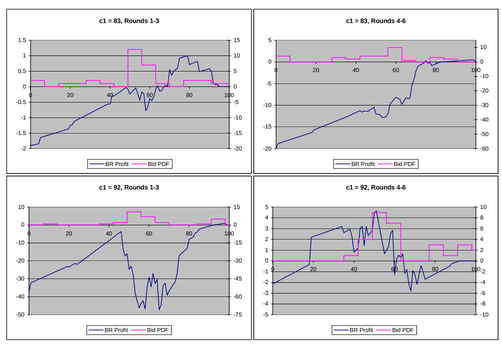

Figure 5: Best Response Profits and pdf of bids in 1U in the first and last three rounds for cost types 83 and 92.

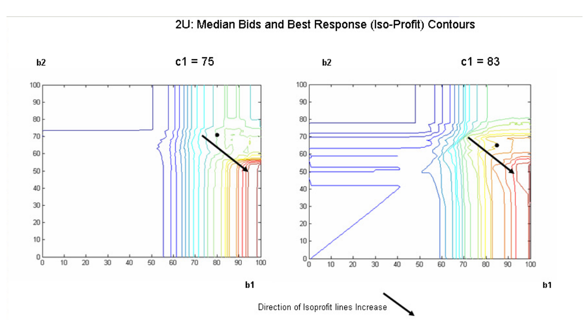

Figure 6: Median Bids and Best Response iso-profit contours

Online Appendix A: Derivation of Equilibrium for 2D and 2U

2D

Consider a slight variation of the rules in which we impose the additional constraint that a winning allocation cannot involve one unit of production from each of the three bidders (so the allocation can only involve one bidder supplying two units and another supplying one unit). Call this variation 2D'. We claim that when cost types are uniformly distributed on $[0, 50]$, it is an equilibrium of 2D' for each bidder to use the following bidding strategies for one and two units respectively:

$$\begin{aligned} b_1^*(c_1) &= \frac{100}{3} + \frac{2}{3}c_1 \\ b_2^*(c_2) &= 25 + \frac{2}{3}c_2 \end{aligned}$$

Proof

Rewrite the bidding functions in the equivalent form $b_1^*(\theta) = 100 - \frac{2}{3}\theta$ and $b_2^*(\theta) = 100\left(\frac{7}{12}\right) + \frac{1}{3}\theta$. Suppose that Suppliers 2 and 3 are using these bidding functions. It suffices to show that it is a best response for Supplier 1 to use these bidding functions as well. As shorthand, write $\beta_i = b_1^*(\theta_i)$ and $B_i = b_2^*(\theta_i)$ for the realizations of Supplier 1's four opposing bids (for $i = 2, 3$). (So β_2 and β_3 are the one unit bids and B_2 and B_3 are the two unit bids.) Note that with $b_1^*(\theta)$ decreasing in θ and $b_2^*(\theta)$ increasing in θ, we have $(B_3 - B_2)(\beta_3 - \beta_2) < 0$ with probability 1. That is, neither Supplier 2 nor Supplier 3 has the both the lowest one unit bid and the lowest two unit bid, among the two of them. Labeling the bids of the supplier drawing $\theta_H = \max(\theta_2, \theta_3)$ with an H and the bids of the supplier drawing $\theta_L = \min(\theta_2, \theta_3)$ with an L, we have $\beta_H < \beta_L$ and $B_L < B_H$. Write β_1 and B_1 for Supplier 1's bids. Under the two supplier constraint, the lowest total cost in an allocation excluding Supplier 1 is

$$\beta_H + 2B_L$$

(one unit from Supplier 2 and two units from Supplier 3). The lowest total cost in an allocation in which Supplier 1 provides one unit is

$$\beta_1 + 2B_L$$

while the lowest total cost in an allocation in which Supplier 1 provides two units is

$$\beta_H + 2B_1$$

Thus, given β_1 and B_1, Supplier 1 provides one unit in the event that

$$\beta_1 + 2B_L < \beta_H + 2B_L \quad \text{and } \beta_1 + 2B_L < \beta_H + 2B_1$$

or equivalently,

$$\beta_H > \beta_1 \quad \text{and} \quad \beta_H - 2B_L > \beta_1 - 2B_1 \qquad (*)$$

Meanwhile, Supplier 1 provides two units in the event that

$$\beta_H + 2B_1 < \beta_H + 2B_L \quad \text{and} \quad \beta_H + 2B_1 < \beta_1 + 2B_L$$

or equivalently,

$$B_L > B_1 \quad \text{and} \quad \beta_H - 2B_L < \beta_1 - 2B_1 \qquad (**)$$

We can write Supplier 1's expected payoff when he bids (β_1, B_1) with type θ_1 as

$$\begin{aligned}
\pi_1(\beta_1, B_1; \theta_1) &= P_1(\beta_1, B_1)(\beta_1 - c_1(\theta_1)) + P_2(\beta_1, B_1)(2B_1 - 2c_2(\theta_1)) \\
&= P_1(\beta_1, B_1)(\beta_1 - 100 + \theta_1) + P_2(\beta_1, B_1)(2B_1 - 100 - \theta_1)
\end{aligned}$$

where

$$\begin{aligned}
P_1(\beta_1, B_1) &= \Pr(\beta_H > \beta_1 \,\cap\, \beta_H - 2B_L > \beta_1 - 2B_1) \quad , \text{ and} \\
P_2(\beta_1, B_1) &= \Pr(B_L > B_1 \,\cap\, \beta_H - 2B_L < \beta_1 - 2B_1)
\end{aligned}$$

The second inequality in each of the probability terms immediately above reflects the fact that Supplier 1 could fail to win one unit because his own two unit bid is too competitive, or *vice versa*. Consider the mathematical expression we would get if we simply left these second inequality terms out:

$$\tilde{\pi}_1(\beta_1, B_1; \theta_1) \equiv \Pr(\beta_H > \beta_1)(\beta_1 - 100 + \theta_1) + \Pr(B_L > B_1)(2B_1 - 100 - \theta_1)$$

Note that $\tilde{\pi}_1(\beta_1, B_1; \theta_1) \geq \pi_1(\beta_1, B_1; \theta_1)$ by construction, since each of the probability terms is less constrained (and hence weakly larger) than the one it replaces.

From here on out, our strategy is as follows. First we solve for the bids (β_1^*, B_1^*) that maximize $\tilde{\pi}_1(\beta_1, B_1; \theta_1)$ and show that these bids correspond to the equilibrium bids $b_1^*(\theta_1)$ and $b_2^*(\theta_1)$. Then we show that at (β_1^*, B_1^*) the values of the true expected payoff $\pi_1(\beta_1^*, B_1^*; \theta_1)$ and our artificial function $\tilde{\pi}_1(\beta_1^*, B_1^*; \theta_1)$ are equal. Thus, since $\tilde{\pi}_1(\beta_1, B_1; \theta_1) \geq \pi_1(\beta_1, B_1; \theta_1)$, (β_1^*, B_1^*) *a fortiori* also maximizes the true expected payoff function $\pi_1(\beta_1, B_1; \theta_1)$.

Step 1: Maximize $\tilde{\pi}_1(\beta_1, B_1; \theta_1)$

Note that

$$\begin{aligned}
\Pr(\beta_H > \beta_1) &= \Pr\left(100 - \frac{2}{3}\max(\theta_2, \theta_3) > \beta_1\right) \\
&= \Pr\left(\max(\theta_2, \theta_3) < \frac{3}{2}(100 - \beta_1)\right) \\
&= \left(\frac{3}{100}(100 - \beta_1)\right)^2 \qquad \text{for } \beta_1 \in \left[66\frac{2}{3}, 100\right]
\end{aligned}$$

given θ_2 and θ_3 uniform and independent on $[0, 50]$. (For β_1 above or below this range, $\beta_H > \beta_1$ is satisfied either never or always.) Similarly, for the other probability term we have

$$\Pr\left(B_L > B_1\right) = \left(\frac{3}{50}\left(75 - B_1\right)\right)^2 \qquad \text{for } B_1 \in \left[58\frac{1}{3}, 75\right]$$

Thus, we must solve

$$\left(\beta_1^*, B_1^*\right) = \arg\max_{(\beta_1, B_1)} \left(\left(\frac{3}{100}\left(100 - \beta_1\right)\right)^2 \left(\beta_1 - 100 + \theta_1\right)\right) + \left(\left(\frac{3}{50}\left(75 - B_1\right)\right)^2 \left(2B_1 - 100 - \theta_1\right)\right)$$

But notice that this maximization can be separated into two pieces: the first half depends only on β_1 and θ_1, while the second have depends only on B_1 and θ_1. Therefore, we have

$$\begin{aligned} \beta_1^* &= \arg\max_{\beta_1} \left(\frac{3}{100}\left(100 - \beta_1\right)\right)^2 \left(\beta_1 - 100 + \theta_1\right) \quad \text{and} \\ B_1^* &= \arg\max_{B_1} \left(\frac{3}{50}\left(75 - B_1\right)\right)^2 \left(2B_1 - 100 - \theta_1\right) \end{aligned}$$

Solving these to show that $\beta_1^* = 100 - \frac{2}{3}\theta_1 = b_1^*\left(\theta_1\right)$ and $B_1^* = 58\frac{1}{3} + \frac{1}{3}\theta_1 = b_2^*\left(\theta_1\right)$ is a standard exercise.

<u>Step 2:</u> $\tilde{\pi}_1\left(\beta_1^*, B_1^*; \theta_1\right) = \pi_1\left(\beta_1^*, B_1^*; \theta_1\right)$

Note that if all three sets of bids are formulated according to the equilibrium bidding functions, then the lowest one unit bid corresponds to the highest draw of θ and the highest two unit bid, and conversely. With this in mind, reconsider $P_1\left(\beta_1^*, B_1^*\right) = \Pr\left(\beta_H > \beta_1^* \,\cap\, \beta_H > \beta_1^* - 2\left(B_1^* - B_L\right)\right)$. If $\beta_H > \beta_1^*$ is satisfied, then $\theta_1 = \max\left(\theta_1, \theta_2, \theta_3\right)$, and therefore, $B_1^* > B_H > B_L$. But of course this means that $\beta_1^* > \beta_1^* - 2\left(B_1^* - B_L\right)$. We conclude that $\beta_H > \beta_1^* \Rightarrow \beta_H > \beta_1^* - 2\left(B_1^* - B_L\right)$, and therefore, that $P_1\left(\beta_1^*, B_1^*\right) = \Pr\left(\beta_H > \beta_1^*\right)$.

Similarly, consider $P_2\left(\beta_1^*, B_1^*\right) = \Pr\left(B_L > B_1^* \,\cap\, 2B_L > 2B_1^* - \left(\beta_1^* - \beta_H\right)\right)$. In this case, if $B_L > B_1^*$ is satisfied, then $\theta_1 = \min\left(\theta_1, \theta_2, \theta_3\right)$ and therefore $\beta_1^* > \beta_L > \beta_H$. But this would imply that $2B_1^* > 2B_1^* - \left(\beta_1^* - \beta_H\right)$. In this case, we conclude that $B_L > B_1^* \Rightarrow 2B_L > 2B_1^* - \left(\beta_1^* - \beta_H\right)$, and therefore that $P_2\left(\beta_1^*, B_1^*\right) = \Pr\left(B_L > B_1^*\right)$.

Together, these suffice to show that $\pi_1\left(\beta_1^*, B_1^*; \theta_1\right) = \tilde{\pi}_1\left(\beta_1^*, B_1^*; \theta_1\right)$. Since $\pi_1\left(\beta_1, B_1; \theta_1\right) \leq \tilde{\pi}_1\left(\beta_1, B_1; \theta_1\right)$ by construction, $\pi_1\left(\beta_1^*, B_1^*; \theta_1\right) = \tilde{\pi}_1\left(\beta_1^*, B_1^*; \theta_1\right)$, and $\tilde{\pi}_1\left(\beta_1, B_1; \theta_1\right)$ attains its maximum value at $\left(\beta_1^*, B_1^*\right)$, we conclude, *a fortiori* , that $\left(\beta_1^*, B_1^*\right)$ maximizes $\pi_1\left(\beta_1, B_1; \theta_1\right)$. Thus, if Suppliers 2 and 3 use the conjectured equilibrium strategy, we have shown that it is a best response for Supplier 1 to use the same bidding strategy, which is what we set out to prove.

Finally, we claim that we do not err too much by using equilibrium bids from 2D' as a benchmark, even though they are not the correct equilibrium bids when three bidder allocations are permitted.

If bidders use the 2D' equilibrium strategies in format 2D, then the rules will select a three bidder allocation only if $\beta_1 + \beta_2 + \beta_3 = 300 - \frac{2}{3}(\theta_1 + \theta_2 + \theta_3)$ is smaller than the least cost two bidder allocation. That least cost two bidder allocation involves buying one unit at $b_1^*(\theta_{\max})$ and two units at $b_2^*(\theta_{\min})$ (where $\theta_{\max} = \max(\theta_1, \theta_2, \theta_3)$ and $\theta_{\min} = \min(\theta_1, \theta_2, \theta_3)$), at a total cost of

$$b_1^*(\theta_{\max}) + 2b_2^*(\theta_{\min}) = 100\left(2\frac{1}{6}\right) + \frac{2}{3}(\theta_{\min} - \theta_{\max})$$

So there is a three bidder allocation whenever

$$300 - \frac{2}{3}(\theta_1 + \theta_2 + \theta_3) < 100\left(2\frac{1}{6}\right) + \frac{2}{3}(\theta_{\min} - \theta_{\max})$$

or equivalently,

$$2\theta_{\min} + \theta_{med} > 125$$

where θ_{med} is defined to be the median of θ_1, θ_2, and θ_3. One can check that with the types uniformly distributed on $[0, 50]$,

$$\Pr(2\theta_{\min} + \theta_{med} > 125) < 0.007$$

so the frequency of a three bidder allocation is quite rare and has a very small impact on the bidders' expected profits. With a bit more work, one could show that the 2D' equilibrium strategies represent an $\varepsilon-$ equilibrium of 2D.

2U

We claim that the following bidding strategy represents a symmetric pure strategy equilibrium when cost types θ are distributed uniformly and independently on $[0, 50]$.

$$\begin{aligned} b_1^*(c_1) &= \begin{cases} 3c_1 + 100\ln\left(1 - \frac{c_1}{100}\right) + 100\ln 3 - 100 & \text{if } c_1 \leq \frac{200}{3} \approx 66.7 \\ 100 & c_1 > \frac{200}{3} \end{cases} \\ b_2^*(c_1) &= 0 \end{aligned}$$

<u>Part 1:</u> Optimality of $b_2^*(\cdot)$

Suppose that Bidders 2 and 3 are playing the equilibrium strategies, and consider the best response of Bidder 1. Suppose that Bidder 1 has provisionally chosen some function b_{11} for its one unit bit and is considering his two unit bid b_{12}. Regardless of how 1 bids, the rationing rules and the bids of 2 and 3 imply that a two unit bid of zero from *some* bidder will always form part of the market-clearing allocation. Thus, any bid $b_{12} > 0$ by 1 has no chance to be accepted. If Bidder 1 chooses $b_{12} > 0$, then

- it never wins two units,

- it wins one unit if $b_{11} < \min(b_{21}, b_{31})$ – call this event (i)

- it wins zero units if $b_{11} > \min(b_{21}, b_{31})$ – call this event (ii)

Consider a switch by Bidder 1 to bidding $b_{12} = 0$. This has no effect on the allocation in event (i) because Bidder 1's first unit bid is needed to achieve the market-clearing price of $p^* = b_{11}$ (achieved by accepting b_{11} and either $b_{22} = 0$ or $b_{32} = 0$). In event (ii), one of the other bidders, say Bidder 2, has the minimal one unit bid, and the price will be $p^* = b_{21}$. A switch to $b_{12} = 0$ gives Bidder 1 a 50-50 chance (with Bidder 3) of supplying the inframarginal two units at a price of b_{21} and earning an additional $2(b_{21} - c_{12})$. But this expression is always strictly positive – the function b_1^* is increasing and so $b_{21} \geq \underline{b}_1 \equiv b_1^*(50) \approx 90.54 > 75 \geq c_{12}$ – so switching to $b_{12} = 0$ at least weakly benefits bidder 1, *regardless of his one unit bidding strategy* b_{11}.

Part 2: Optimality of $b_1^*(\cdot)$

Now fix the equilibrium strategies for Bidders 2 and 3 and fix $b_{12} = b_2^*$. We need to show that $b_{11} = b_1^*$ is a best response. The price and allocation will have the form $p^* = \min\{b_{11}, b_{21}, b_{31}\}$, with Bidder i supplying 1 unit, and earning $b_{i1} - c_{i1}$, if $b_{i1} = p^*$ and supplying 2 units with probability $\frac{1}{2}$ if $b_{i1} > p^*$, earning $2(p^* - c_{i2})$. That is, the most competitive 1 unit bidder supplies the marginal unit and sets the price, while the other two bidders, having tied at a two unit bid of zero, have equal chances of supplying those two inframarginal units. Define the variable $x = \min\{c_{21}, c_{31}\}$. Then (because b_1^* is increasing) Bidder 1 earns $b_{11} - c_{11}$ if $b_{11} < b_1^*(x)$ and (in expectation) earns $b_1^*(x) - c_{12}$ if $b_{11} > b_1^*(x)$. These inequalities are equivalent to $b_1^{*-1}(b_{11}) \lesseqgtr x$. Since x has a cumulative distribution function given by $G(x) = 1 - (1 - F(x))^2$, where $F(x) = \frac{x-50}{50}$ is the c.d.f. of $c_1\tilde{}U(50, 100)$, we can write the expected profit to Bidder 1 from a bid $b_{11} = b$ as

$$\begin{aligned}
\pi_1(b) &= \left(1 - G\left(b_1^{*-1}(b)\right)\right)(b - c_{11}) + \int_{50}^{b_1^{*-1}(b)} g(x)\left(b_1^*(x) - c_{12}\right) dx \\
&= \left(1 - F\left(b_1^{*-1}(b)\right)\right)^2 (b - c_{11}) + 2\int_{50}^{b_1^{*-1}(b)} f(x)(1 - F(x))\left(b_1^*(x) - c_{12}\right) dx
\end{aligned}$$

Taking first order conditions and simplifying, we have

$$\frac{d}{db}(\pi_1(b)) = \left(1 - F\left(b_1^{*-1}(b)\right)\right)^2 + 2\frac{f\left(b_1^{*-1}(b)\right)\left(1 - F\left(b_1^{*-1}(b)\right)\right)}{b_1^{*\prime}\left(b_1^{*-1}(b)\right)}(c_{11} - c_{12})$$

Recall that $c_{12} = 100 - c_{11}/2$, so for $c_{11} > \frac{2}{3}100$, we have $c_{11} > c_{12}$, in which case both terms in this expression are strictly positive. It follows that setting $b_{11} = 100$ (the highest permissible bid) maximizes $\pi_1(b)$ if $c_{11} > \frac{2}{3}100$. For $c_{11} < \frac{2}{3}100$, we solve the FOC $\frac{d}{db}(\pi_1(b)) = 0$ to get the hazard rate condition:[25]

$$\frac{1 - F\left(b_1^{*-1}(b)\right)}{f\left(b_1^{*-1}(b)\right)} = 2\frac{c_{12} - c_{11}}{b_1^{*\prime}\left(b_1^{*-1}(b)\right)} = \frac{200 - 3c_{11}}{b_1^{*\prime}(x_b)}$$

Substituting in for F and $b_1^{*\prime}(c) = \frac{200-3c}{100-c}$, the optimal b satisfies

$$\begin{aligned} 100 - b_1^{*-1}(b) &= \frac{200 - 3c_{11}}{\frac{200-3c_{11}}{100-c_{11}}} = 100 - c_{11} \quad , \quad \text{or} \\ b_1^{*-1}(b) &= c_{11} \quad \text{, and therefore,} \\ b &= b_1^*(c_{11}) \end{aligned}$$

In Parts 1 and 2, we optimize one component of Bidder 1's strategy while holding another component fixed. By itself, this does not suffice to show that (b_1^*, b_2^*) is a best response for Bidder 1. To show that (b_1^*, b_2^*) does weakly better for Bidder 1 than arbitrary alternative bidding functions $\left(\tilde{b}_1, \tilde{b}_2\right)$, consider switching from $\left(\tilde{b}_1, \tilde{b}_2\right)$ to (b_1^*, b_2^*) in two steps:

$$\begin{aligned} \left(\tilde{b}_1, \tilde{b}_2\right) &\to \left(\tilde{b}_1, b_2^*\right) \\ \left(\tilde{b}_1, b_2^*\right) &\to (b_1^*, b_2^*) \end{aligned}$$

Bidder 1's expected payoff weakly improves at the first step by Part 1 and weakly improves at the second step by Part 2, so (b_1^*, b_2^*) is a best response for Bidder 1 as claimed.

(Second order condition for optimality of b_1^*)

We are concerned with the case in which $c_{11} - c_{12} < 0$. Define $x = b_1^{*-1}(b)$ and $\Delta = c_{12} - c_{11} > 0$ so that we have

$$\begin{aligned} \frac{d}{db}(\pi_1(b)) &= (1 - F(x))^2 - 2\frac{f(x)(1 - F(x))}{b_1^{*\prime}(x)}\Delta \\ &= f(x)(1 - F(x))\left(\frac{1 - F(x)}{f(x)} - \frac{2\Delta}{b_1^{*\prime}(x)}\right) \end{aligned}$$

The second derivative is then

$$\begin{aligned} \frac{d^2}{db^2}(\pi_1(b)) &= \frac{dx}{db} \cdot \frac{d}{dx}\left(f(x)(1 - F(x))\left(\frac{1 - F(x)}{f(x)} - \frac{2\Delta}{b_1^{*\prime}(x)}\right)\right) \\ &= \frac{dx}{db} \cdot \left(\left(\left(\frac{1 - F(x)}{f(x)} - \frac{2\Delta}{b_1^{*\prime}(x)}\right)\frac{d}{dx}(f(x)(1 - F(x)))\right) + \left(f(x)(1 - F(x))\frac{d}{dx}\left(\frac{1 - F(x)}{f(x)} - \right.\right.\right. \end{aligned}$$

[25] See below for the second order condition. We may be able to omit this.

Notice that $\left(\frac{1-F(x)}{f(x)}-\frac{2\Delta}{b_1^{*\prime}(x)}\right)$ is just the FOC, so the first term above will drop out when we evaluate this expression at the optimal bid. For the second term, note that $f'(x)=0$, so $\frac{d}{dx}\left(\frac{1-F(x)}{f(x)}\right)=-1$, and by direct computation, $b_1^{*\prime\prime}(x)=-\frac{100}{(100-x)^2}<0$. Furthermore, $\frac{dx}{db}>0$ because b_1^* is increasing. Finally, evaluated at the optimal bid, we have $x=c_{11}$. In summary, we have

$$\begin{aligned}\frac{d^2}{db^2}\left(\pi_1(b)\right)|_{b=b_1^*(c_{11})} &= \left(\frac{dx}{db}\right)\left(\left(f\left(c_{11}\right)\left(1-F\left(c_{11}\right)\right)\right)\right)\left(-1+\frac{2\Delta}{\left(b_1^{*\prime}\left(c_{11}\right)\right)^2}b_1^{*\prime\prime}\left(c_{11}\right)\right)\\ &< 0\end{aligned}$$

since the first two terms are positive and the third is negative.

Online Appendix B: Statistical Analysis of Bidding Behavior

This appendix complements the regression results discussed in the main text with somewhat more complete discussion of robustness and alternative specifications.

Clustering at the sub-session level

Because the same subjects participate in many auctions within a single treatment, independence is of some concern. A typical approach to this issue would be to treat the data as a panel with subject fixed effects, but here there are reasons to take a different approach. First of all, our objective is a cross-treatment comparison of outcomes (e.g., auction efficiency) that are *jointly* determined, in a non-linear way, by the actions of several subjects simultaneously. A consistent, idiosyncratic bidding pattern by a particular subject can, and typically will, have very different implications for auction efficiency depending on which opponents she is matched with in any given round.[26] Furthermore, subjects change costs every six rounds – we have no basis for thinking that a subject's tendencies with a low c_1 type predict her bidding patterns with a high c_1 type, or *vice versa*. However, a stronger case can be made for correlation at the level of the sub-session. Patterns in auction outcomes that emerge jointly out of a particular assignment of individuals to cost types will be common to all auctions within a sub-session and less correlated with outcomes across sessions, after cost types have been randomly reassigned. For this reason, all results the results we present have been computed with robust standard errors corrected for clustering at the sub-session level.

Alternative specifications: Allowing coefficients to vary across auction formats

These basic specifications restrict the differences among treatments to a single difference in levels, but we might also suspect that the auction formats respond differentially to different realizations of supplier costs. One crude but succinct way to summarize these cost realizations is by using the median supplier type. Small values of *c1med* are correlated with realizations in which the three suppliers tend to have increasing average costs, while large values of *c1med* are correlated with realizations in which the suppliers tend to have decreasing average costs. To simplify the discussion below, we will say that suppliers tend to be small when $c1med$ is small, when what we mean is that the tendency of suppliers' average costs to be increasing or decreasing is statistically related to $c1med$ in the manner described above. Similarly, we will say that suppliers tend to be large when $c1med$ is large. Below we present sample averages for *Efficiency*, *Markup*, and *Buyercost* conditioned on whether *c1med* is small or large.

[26] In principle, one could associate a fixed effect with every possible triple of subjects, but vastly more data would be required to do this.

	Efficient		*Markup*		*BC*	
$c1med$	< 75	≥ 75	< 75	≥ 75	< 75	≥ 75
$1U$	.45	.25	12.8	-8.4	194.2	182.3
$2U$	.49	.81	22.0	49.1	196.3	222.3
$2D$	.77	.76	18.3	15.7	186.4	201.2

Figure 7 presents a slightly different slice at the same data. Rather than pooling the values of $c1med$ into two groups (< 75 and ≥ 75), for each auction format we plot the average for each outcome variable ($Efficiency$, $Markup$, and *Buyercost*) at each value of $c1med$ (50 through 100). (So each plot has three lines, one for each treatment.) These plots and the table above reveal some interesting patterns. Broadly, the two uniform price formats do not differ too much when the suppliers tend to have increasing average costs ($c1med < 75$), but they diverge dramatically when supplier costs tend to be decreasing. As we shift from $c1med < 75$ to $c1med \geq 75$, efficiency, profit margins, and total procurement cost all fall under one-part bidding, while all three rise under two-part bidding. In contrast, shifting the mix of suppliers from smaller types toward larger types has a fairly muted impact under format $2D$ (although procurement cost does rise as more suppliers exhibit decreasing average costs). The factors contributing to inefficient allocations under one-part bidding with avoidable costs are familiar from other work in the auction literature. However, some of the other patterns are less familiar and perhaps puzzling. In particular, the fact that the relative attractiveness of one-part pricing to the buyer (that is, *BC*) *improves* as suppliers types grow larger, despite ballooning inefficiency in production, is rather striking. Before undertaking an explanation, we first return to the regression analysis begun above in order to put these summary statistics on a more rigorous footing.

For each of the outcome measures considered earlier, we repeat our earlier analysis with a full set of interaction effects that permits all of the right-hand side coefficients to vary across treatments; these results are presented in Column 3. Most of these differences are not significant, and the specifications in Column 2 focus on one set of interactions that is consistently highly significant: the treatment-specific effects of *c1med*. Note that the regressions control for the cost types of the highest and lowest type supplier in an auction, so the coefficient on *c1med* isolates the effect of shifting the median supplier toward a higher c_1 type. As Column 2 shows, this median supplier type plays quite different roles in the three auction formats. In format $1U$, shifting the median supplier toward higher types reduces the chance of an efficient allocation, but also reduces profit margins. The impact of the latter change is greater, and so total buyer cost goes down. Under format $2D$, the composite effects (the sum of the coefficients on *c1med* and *c1med*·$2D$) are significant[27]

[27] Only at the 10% level for the *Efficiency* results.

and have the same sign as for $1U$, but the magnitude of all of the effects is substantially smaller. However, for $2U$, all three effects are reversed: an increase in *c1med* improves efficiency, raises profit margins, and increases the overall procurement cost to the buyer. For concreteness, the table below summarizes these effects for an example in which the median supplier is shifted by one type – from constant marginal cost ($(c_1, c_2) = (67, 67)$) to slightly decreasing average costs ($(c_1, c_2) = (75, 63)$).[28]

Effect of Shifting the Median Supplier Toward Decreasing Average Costs

	$\Delta \Pr(Efficient)$	$\Delta Markup$	ΔBC
$1U$	-0.10	-13.2	-9.1
$2U$	0.06	5.3	4.7
$2D$	-0.04	-4.1	-2.2

The regression results are broadly consistent with the summary statistics reported in section 5. We will provide only a brief heuristic explanation of these results here; a more thorough analysis is undertaken in the main text. Inefficiency arises under $1U$ in large part because one-dimensional bid functions can do no better than to achieve a fixed ordering of cost types, from inframarginal to marginal to out-of-the-market. However, efficiency may require the dispatch order of two suppliers to depend on the particular cost realization of the final supplier, which is impossible to achieve. However, profit margins and buyer cost will depend on how suppliers strategically try to maneuver into their preferred positions in the (uncertain) dispatch order. Broadly speaking, suppliers with decreasing average costs face the danger of being rationed to one unit at an unprofitable price as the marginal bidder. They can try to avoid this outcome either by bidding more competitively (to be more likely to be inframarginal) or by bidding less competitively (to be sure to cover one unit cost). Our subjects predominantly take the former approach, and as a result, the market-clearing price grows very competitive as the market composition shifts toward higher c_1 types.

Two-part bidding under format $2U$ effectively creates competition in two sub-markets – the “market” to supply two units and the “market” to supply one unit.[29] The lowest price that can procure from both sub-markets tends to be set by bidding in the sub-market with higher bids and weaker competition, and for our parameters, this is usually the one unit sub-market. Thus, shifting the median supplier toward higher c_1 types tends to shift competition away from the sub-market where it would be more likely to have an effect on the price, resulting in a higher procurement cost for the buyer. As with $2U$, two-part bidding in $2D$ effectively creates submarkets for one unit and two

[28] Probability changes for the probit are evaluated starting from the mean probability of efficiency for each treatment.

[29] A three bidder allocation is possible in principle, but occurs once in 120 auctions in the data.

units, but under $2D$ each of these sub-markets gets its own price. As a result, shifting competition from the one unit to the two unit sub-market (*via* an increase in $c1med$) has only a minor impact on the *total* procurement cost.

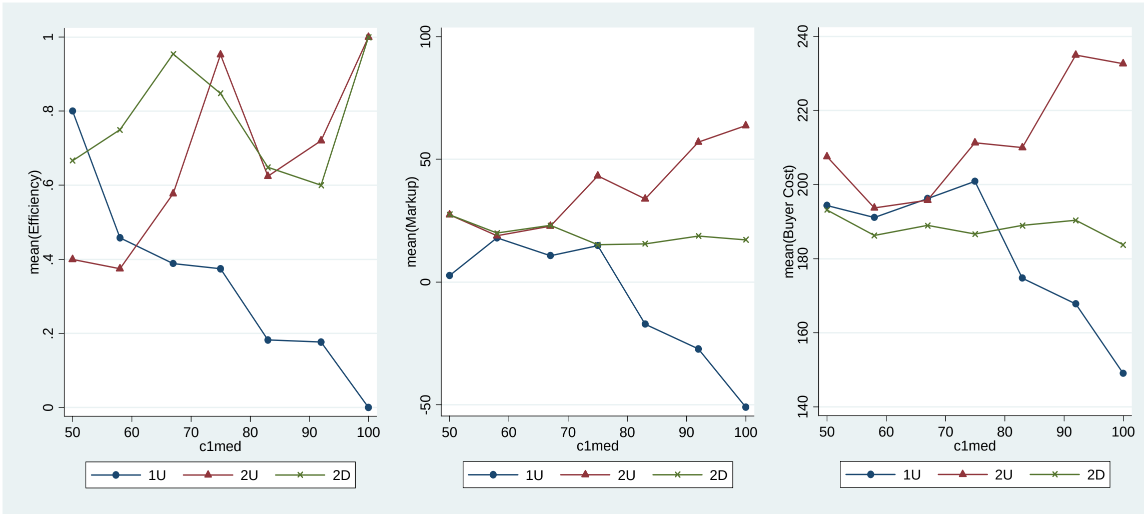

Figure 7: Mean auction outcomes vs. c1med, by auction format

Online Appendix C: Tables

Probability of Efficient Allocation (Probit)

Efficient	(1)	**(2)**	(3)
period	0.000	**0.001**	-0.002
	(0.03)	**(0.30)**	(0.33)
modperiod	-0.021	**-0.018**	-0.062
	(0.58)	**(0.49)**	(1.88)*
c1min	0.003	**0.001**	0.004
	(0.58)	**(0.20)**	(0.49)
c1med	-0.009	**-0.038**	-0.041
	(0.94)	**(3.53)*** **	(3.46)***
c2min	-0.020	**-0.020**	-0.027
	(1.98)**	**(2.04)** **	(1.81)*
u2	0.928	**-3.695**	-5.021
	(6.11)***	**(4.71)*** **	(3.00)***
d2	1.186	**-0.398**	-1.314
	(8.46)***	**(0.44)**	(0.72)
c1med_u2		**0.062**	0.071
		(5.42)* **	(4.82)***
c1med_d2		**0.022**	0.025
		**(1.72)* **	(1.78)*
c1min_u2			-0.010
			(0.70)
c1min_d2			-0.000
			(0.02)
c2min_u2			0.017
			(0.76)
c2min_d2			0.007
			(0.25)
period_u2			0.001
			(0.09)
period_d2			0.009
			(0.87)
Modperiod_u2			0.086
			(0.89)
Modperiod_d2			0.063
			(0.91)
Constant	1.187	**3.405**	4.027
	(1.26)	**(3.60)*** **	(3.35)***
Observations	420	**420**	420
Log pseudolikelihood	-249.533	**-237.408**	-236.648
Pseudo R^2	0.122	**0.165**	0.167

Robust z statistics in parentheses

- significant at 10%; ** significant at 5%; *** significant at 1%

Supplier Profit Margin (OLS)

Markup	(1)	**(2)**	(3)
period	0.121	**0.155**	-0.010
	(0.74)	**(1.09)**	(0.03)
modperiod	0.083	**0.203**	-1.824
	(0.08)	**(0.20)**	(0.78)
c1min	0.187	**0.122**	-0.297
	(0.98)	**(0.67)**	(1.45)
c1med	-0.466	**-1.645**	-1.502
	(1.89)*	**(10.12)***	(9.18)***
c2min	-1.430	**-1.472**	-1.660
	(5.73)***	**(5.40)***	(3.44)***
u2	38.413	**-135.455**	-186.996
	(9.05)***	**(9.53)***	(3.25)***
d2	16.802	**-66.853**	-125.047
	(4.00)***	**(5.82)***	(2.55)**
c1med_u2		**2.311**	1.816
		(11.73)*	(6.72)***
c1med_d2		**1.135**	1.231
		(7.01)*	(6.38)***
c1min_u2			1.099
			(4.08)***
c1min_d2			0.066
			(0.25)
c2min_u2			0.138
			(0.18)
c2min_d2			0.588
			(1.02)
period_u2			0.356
			(0.98)
period_d2			0.275
			(0.79)
modperiod_u2			2.801
			(1.12)
modperiod_d2			2.995
			(1.23)
Constant	100.317	**192.427**	226.931
	(2.98)**	**(8.15)***	(5.51)***
Observations	420	**420**	420
R-squared	0.25	**0.39**	0.41

Robust t statistics in parentheses
* significant at 10%; ** significant at 5%; *** significant at 1%

Total Procurement Cost to Buyer (OLS)

Buyercost	(1)	**(2)**	(3)
period	0.113	**0.138**	-0.006
	(1.09)	**(1.37)**	(0.04)
modperiod	-0.102	**-0.013**	-2.247
	(0.12)	**(0.02)**	(1.30)
c1min	0.979	**0.931**	0.194
	(4.24)***	**(4.20)*** **	(1.01)
c1med	-0.251	**-1.133**	-0.775
	(1.66)	**(6.57)*** **	(5.43)***
c2min	0.444	**0.413**	0.550
	(1.90)*	**(1.63)**	(1.05)
u2	24.813	**-104.690**	-139.854
	(8.75)***	**(8.06)*** **	(2.74)**
d2	1.826	**-61.320**	-90.784
	(0.63)	**(5.65)*** **	(2.29)**
c1med_u2		**1.721**	0.889
		(8.89)* **	(3.65)***
c1med_d2		**0.856**	0.656
		(5.28)* **	(4.21)***
c1min_u2			1.612
			(6.59)***
c1min_d2			0.465
			(1.91)*
c2min_u2			-0.308
			(0.43)
c2min_d2			0.041
			(0.08)
period_u2			0.386
			(1.80)*
period_d2			0.228
			(1.09)
modperiod_u2			3.376
			(1.81)*
modperiod_d2			2.989
			(1.66)
Constant	121.745	**190.600**	210.783
	(4.78)***	**(11.83)*** **	(5.50)***
Observations	420	**420**	420
R-squared	0.28	**0.38**	0.45

Robust t statistics in parentheses
* significant at 10%; ** significant at 5%; *** significant at 1%