Procurement auctions with avoidable fixed costs: an experimental approach

Larson, Nathan and Elmaghraby, Wedad

University of Virginia

2008

Procurement Auctions with Avoidable Fixed Costs: An Experimental Approach

Wedad Elmaghraby[*] and Nathan Larson[†‡]

January 11, 2011

Abstract

Bidders in procurement auctions often face avoidable fixed costs. This can make bidding decisions complex and risky, and market outcomes volatile. If bidders deviate from risk neutral best responses, either due to faulty optimization or risk attitudes, then equilibrium predictions can perform poorly. In this paper, we confront laboratory bidders with three auction formats that make bidding difficult and risky in different ways. We find that measures of 'difficulty' provide a consistent explanation of deviations from best response bidding across the three formats. In contrast, risk and loss preferences cannot explain behavior across all three formats.

Keywords: Auctions, Experimental, Procurement, Synergies, Asymmetric Bidders, Learning, Optimization errors

1 Introduction

Procurement auctions are often used in settings where suppliers have avoidable fixed costs (Elmaghraby (2007)). Prior experimental studies have showed that bidding and auction outcomes in such settings can be quite volatile, and that convergence to equilibrium predictions may be poor. In this paper, we use experiments to study the reasons that suppliers deviate from risk neutral best response bidding when avoidable fixed costs are present. In particular, we ask whether these deviations are better explained by preferences (e.g., suppliers who are not risk neutral) or by optimization errors. To address this question, we study bidding in three auction formats across which the level of payoff risk and the difficulty of optimization vary. Our main finding is that measures of optimization difficulty help to explain behavior across auction formats. In contrast, our preference-based explanations do not generalize well across auction formats – while they may be helpful in

understanding bidding under one set of rules, they do not offer reliable guidance about what to expect when the rules change.

Avoidable fixed costs are fixed production costs that a supplier does not incur unless it sells a strictly positive quantity. This makes costs non-convex: the marginal cost of the first unit sold may be higher than that of additional units. One result of this is that small, marginal changes can have large, non-marginal consequences. For example, in an efficient allocation, a small cost rise for one supplier may cause it to be dropped from the allocation completely (rather than having its quantity decline gradually). Furthermore, prior research shows (e.g. Van Boening and Wilcox (1996) and Van Boening and Wilcox (2005)) that even when supplier costs are stable, market prices and allocations can remain volatile rather than settling down toward equilibrium. To some extent, the first point may be directly to blame for the second. Any auction format that aspires to be efficient must be willing to make large changes to the winning allocation in response to small changes in bids. However, this may confront bidders with risks they prefer to avoid or with strategic challenges that confound learning.

We select three auction formats for which we expect the difficulty of optimization and the level of payoff risk to differ. In all three formats, suppliers have (private) average costs that are either increasing, constant, or decreasing, and a supplier can produce either zero, one, or two units. We intepret the decreasing average cost case to represent a supplier with an avoidable fixed cost. For such a supplier, selling at certain prices will be profitable only if she sells her full capacity. In the first auction format $(1U)^1$, each supplier submits a single bid specifying the lowest per-unit price it is willing to accept. Bids are ranked in increasing order, and the auction clears at the lowest uniform price that procures the total desired quantity. The second auction format (2U) retains a uniform price, but allows suppliers to fully express their costs with quantity-dependent bids. In the third format (2D), both bids and prices are quantity dependent. Each format presents some idiosyncratic features. For example, in 1U equilibrium bidding is inherently unstable (in the sense that some bidders must use mixed strategies) and requires some suppliers to bid well below cost and hope for a favorable price and quantity. In 2U, optimal bids are simple but not particularly intuitive – they may elude a subject who is not prepared to experiment with different bids. In 2D, while the winner determination process is complex, equilibrium (for our parameters) tends to segment competition into sub-markets at different quantities. A subject who treats these sub-markets as independent errs only slightly, and this substantially simplifies the strategic problem she faces.

First we test how well equilibrium outcomes (such as efficiency and total procurement cost) predict actual outcomes across these environments. The equilibrium predictions are fairly successful for 2D, but less so for 1U and 2U, suggesting that there are persistent deviations from risk neutral best response bidding. To assess those deviations, we compute empirical best response profits for each subject; these are the profits that a subject would have earned in each auction by best responding to the empirical distribution of bids that he had faced in the past. Since a subject does not face equilibrium bids from opponents, these empirical best response profits provide a

---

[1]Formal descriptions appear in Section 3.3.

more reasonable benchmark than equilibrium payoffs do for the amount that a risk neutral best responder could have earned. In treatment 2D, the actual profits earned by subjects were close to these benchmark profits. However, in the other two environments, actual payoffs are below these benchmarks on average, and do not track them closely from auction to auction.

This leads us to try to develop a unified model of bidding that can explain why bids are close to risk neutral best responses in 2D but deviate from them in the other two treatments. We construct and estimate a learning model in the spirit of Erev and Roth (1998) and Camerer and Ho (1999). In the model, a subject's propensity to pick a particular bid evolves in response to her market experience. Propensities map into probabilities of choosing different bids according to a standard logit formulation. The propensity includes a term for the expected profit that the bid would have earned against past opponents; a subject choosing on this basis alone would play a risk neutral best response. We also allow this propensity to depend on factors that are intended to capture alternative preferences or the difficulty of optimization. We test the robustness, or generality, of these preference and optimization-based explanations by fitting models to two of the three auction formats and then predicting bids in the third format.

We focus on three potential sources of optimization difficulty: volatile feedback, difficulty of the payoff landscape, and non-separability of the bidding problem. Volatile feedback, or 'noisy optimization,' captures the idea that a supplier may have trouble discerning the *expected* payoff of a bid when the bid's realized payoff varies a lot. By 'landscape difficulty' we mean the idea that a payoff function shaped like Mt. Everest may be easier to maximize than one shaped like San Francisco (many local maxima) or like Kansas (long flat stretches). As one might imagine, formalizing this notion poses its own challenges, and we defer discussing our results for this explanation until later. By 'separable,' we mean a bidding problem that can be tackled effectively by decomposition into smaller pieces. We also study a factor that could aid optimization - namely the fact that some strategies are weakly dominated and others are not. For preference-based explanations of bidding, we concentrate on simple formulations of risk and loss aversion.

In some cases we can compare the preference and optimization-based explanations head to head. Both the noisy optimization story and risk averse preferences imply that a bid's payoff variability (defined as the variance of its payoff against a subject's past opponents) should affect how often it is chosen. However because the predicted effects are different, we can compare these explanations in the data. We find that risk aversion cannot provide a consistent explanation of bidding across all three auction formats. In contrast, a single model of noisy optimization explains patterns within and across auction formats rather well.

Next we set weak dominance and loss aversion head to head. All bids below cost, which we call loss exposed, have the potential to lose money. Some of these loss exposed bids are also weakly dominated, while others can be profitable, or even optimal, on average.[2] We test whether (i) subjects avoid all loss exposed bids or (ii) subjects specifically avoid weakly dominated bids. Both (i) and (ii) help to explain bidding, but the predictions of (ii) are more robust across auction

---

[2]This is true for 1U and 2U. For 2D, the loss exposed and weakly dominated sets are the same.

formats. We infer that loss attitudes, if present, are unstable. From (ii) we learn that among bids with low expected payoffs, subjects are better at avoiding the "sure losers" than bids with some chance of an upside.

For both the risk and loss-based models, poor robustness is related to major behavior differences between formats 1U and 2U. In the latter, subjects avoid modestly risky or loss exposed bids, even when the potential payoff gains are large. In the former, subjects favor quite risky bids, even though the payoff gains they offer are small or even negative. Thus, while risk and loss attitudes may be helpful in understanding bidding under one set of rules, they do not offer reliable guidance about what to expect when the rules change.

With regard to separability, in our setup most allocations will involve one two unit supplier and one one unit supplier. It could be tempting for subjects to try to simplify the bidding problem by treating these as separate and independent market niches. To test this, we decompose a bid's expected payoff into portions earned by winning one or two units respectively. Then we enter these two terms separately in the propensity to bid equation. We find that subjects consistently focus on either one unit profits or on two unit profits, ignoring the possibility of gains or losses from winning the other quantity. We argue that this "two separate submarkets" shortcut is closer to the truth in format 2D than in the other auction formats, and this contributes to better profit maximization by subjects in 2D.

Concentrating on expected profits from one submarket is an example of a heuristic that subjects use to simplify bidding. However, even this may be difficult: it requires subjects to understand winner determination well enough to estimate counterfactual market outcomes. Motivated by models of reinforcement learning, we also test for simpler a simpler heuristic: does a subject simply reinforce bids that have won her a particular quantity in the past? We incorporate 'quantity reinforcement' terms into the propensity to bid equation (allowing for spillover reinforcement to nearby bids). These terms let us test for three types of goal-oriented behavior: local inertia (a tendency to stay near any bids used previously), win targeting (favoring any bids that have won positive quantities), and competitive quantity targeting (favoring bids that win the quantity the subject can produce at least average cost). We find that inertia is quite strong in all three auction formats. This serves subjects well in format 2D – their instincts about where to start bidding (cost plus a markup for each quantity) are basically correct and require only local fine-tuning. However, in the other two formats, their initial intuition about how to bid is less accurate, and inertia leads them to overlook opportunities to improve. We also find evidence of quantity targeting, suggesting that subjects use it as a low tech proxy for identifying profitable bids.

In Section 2 we position our contribution relative to related work on markets with non-convex preferences and learning to optimize in complex environments. Section 3 describes our model, auction formats, and experimental procedure. Section 4 presents session-level comparisons of experimental outcomes with theoretical predictions. Section 5 introduces our models of the individual bidding decision, and Section 6 presents results for these models. In Section 7, we test the quantity reinforcement heuristics. We conclude with some suggested directions for future work.[3]

---

[3]Proofs and derivations are in the appendix. Additional information about computational methods is available in

2 Context and Related Work

Avoidable fixed costs are relatively common in procurement; for example, when goods or services are produced to order, there is often a start up cost related to setting up a production line or training employees to handle customized aspects of the order. Wholesale electricity procurement provides a nice illustration of the complications this can cause, and also of competing views about how they should be handled. With certain types of power plants, a start-up cost is incurred whenever an idle plant is called into service. In the original design of the California Power Exchange (one of the earliest deregulated electricity markets in the U.S.), bids allowed suppliers to express variable cost but not to separately express start-up costs. Market clearing involved a simple ranking of bids, with all winning bidders receiving the same uniform price. Because a marginal supplier could be rationed, a supplier faced the risk of not winning a large enough quantity to recoup its start-up cost. In this market, the task of trying to bid competitively enough to win while avoiding this exposure risk was shouldered by the bidders.

In contrast to California, the Pennsylvania-New Jersey-Maryland (PJM) market (the largest competitive wholesale electricity market in the world) allows suppliers to submit multi-part bids, so that it is possible to express both start-up and variable costs. These bids are fed into a complex optimization procedure which determines the winning allocation and sets what amount to bidder-specific prices.[4] This format makes a bidder's decision problem simpler in some respects (it is easy ensure that one's costs are covered) but harder in others: bidding optimally requires an understanding of how one's bid will affect market outcomes (prices and allocations), and the optimization procedure can make this rather opaque. Studying how bidders handle a more opaque allocation procedure without exposure risk, relative to a more transparent allocation where exposure risk is severe, is one object of our comparison between formats 1U and 2U and 2D.

Avoidable fixed costs may be thought of as a type of cost synergy, since a supplier's per unit cost may be lower at larger quantities. There is also a closely related experimental literature on standard (i.e. forward) auctions in which bidders have demand synergies. Katok and Roth (2004) study the decisions of a single large bidder with a demand synergy for two units, who competes against several small (one unit) bidders who follow simple dominant strategies. In an ascending price format where the large bidder faces exposure risk, they find that this bidder bids too cautiously (relative to his risk neutral best response), depressing the seller's revenues (relative to equilibrium). They also study a descending price auction that gives the large bidder more control over its quantity, thus eliminating exposure risk. They find that in this format, the large bidder suffers fewer losses and allocative efficiency is higher. Our experiment also features small and large bidders, with large bidders subject to exposure risk, but there are a number of major differences. Most importantly, we allow for competition among more than one large bidder; this makes it substantially more difficult for bidders to anticipate both how their opponents will bid, and how their own bidding will affect the winning allocation. We also test whether subjects' responses to exposure risk are

---

an online appendix at http://people.virginia.edu/ nl2a/Papers/Submitted/lumpy-code.html .

[4]There is a uniform per-unit price, supplemented by bidder-specific "transfer payments."

consistent across different settings (1U and 2U). Another difference is that our auction formats are static rather than dynamic. Like Katok and Roth, we find more efficient allocations when bidders can make quantity dependent bids (2D and 2U versus 1U), but in our case, fear of exposure does not explain departures from best response bidding very consistently.

Our work is close in spirit to a sequence of three papers by Van Boening and Wilcox (1996 and 2005, henceforth VBW) and Durham et al. (1996). These papers share an experimental design in which multiple buyers (3 or 4) and suppliers (3 or 4) trade in a double auction format. All suppliers have zero variable cost and positive avoidable fixed costs; "larger" suppliers have larger fixed costs and larger capacities. The mix of supplier types is the same in each auction. In VBW (1996), suppliers were restricted to making a single, per unit bid and were paid a uniform price if selected for production. Prior studies have shown this mechanism to be highly efficient when supplier costs are convex. However, with avoidable fixed costs, VBW find a "roller coaster" between efficient and inefficient outcomes that does not settle down as subjects become more experienced. The volatility of outcomes occurs even in parameter configurations for which a competitive equilibrium (at a uniform price) exists. Interesting, they observe a tendency for bids to converge to a single uniform price, even in settings where the only possible competitive equilibrium would require nonlinear prices.

Durham et al. (1996) study whether a richer strategy space for bidders can encourage convergence to efficient outcomes and equilibrium prices in the double auction. They allow suppliers to bid a two-part tariff (a fixed amount if selected to produce, plus a per unit price) and to set an upper bound on the number of units supplied. Despite the expanded bid space, the authors find that the volatility of market outcomes persists and that those outcomes are rarely to never consistent with equilibrium.

VBW (2005) expanded the bidding space further to allow suppliers to offer (a limited number of) price-quantity contracts. (This is their "BUDA" auction format.) One aim was to determine whether suppliers were able to converge to an efficient equilibrium supported by quantity-dependent prices. They found a modest improvement in efficiency (relative to a standard double auction), but persistence of linear (not quantity-dependent) prices. Thus, the evidence about whether the BUDA rules nudge the market toward an equilibrium outcome is mixed. In a variation (RBUDA), one unit contracts were forbidden to see whether the market's tendency toward a (non-equilibrium) uniform price could be broken. (One rationale would be that, by facilitating arbitrage between contracts of different quantities, one unit contracts tend to linearize prices.) VBW found that this restriction does promote convergence to quantity-dependent pricing, but does not improve efficiency (relative to BUDA). Since the equilibrium outcome *is* efficient, it is unclear whether the RBUDA rules pushed the market closer to equilibrium outcomes.

Like these three papers, we study markets with avoidable fixed costs on the supply side, but we study static auctions with a single buyer, not double auctions. We also differ in that our suppliers have a range of convex and non-convex cost structures, and compared to the two VBW papers, bidders are observed over more auction rounds.[5] Like these earlier papers, we also find volatility

[5] Our sessions lasted for 30 rounds, compared with 9 to 16 rounds in the two VBW papers. This does not necessarily

in our market outcomes. However, we diverge from this earlier work by attempting to explain the divergence between actual and equilibrium market outcomes using models of the individual's bidding decision.

While models of bounded rationality and imperfect optimization are common in the economics literature, less attention has been paid to the features of a strategic decision problem that make it hard. Our metrics of difficulty in this paper are partially motivated by a broad literature on this question outside of economics. In the context of a decision task under adverse selection, Bereby-Meyer and Grosskopf (2008) cite extensive evidence that subjects do not learn to avoid bad decisions when payoff variability is high. In a manipulation where payoff variability is reduced, they show that learning improves substantially.

If one treats a payoff function as a 'landscape' traversed by subjects in search of the highest peak, many candidate notions of what constitutes a hard landscape have been proposed, none entirely satisfactory. Jones and Forrest (1995) discuss the challenges of finding a metric that can classify both 'rugged' problems (e.g. San Francisco) and 'needle in a haystack' problems (e.g. a landscape with one ice cube atop an ice rink) as hard. We propose a simple metric based on counting local maxima that can capture both of these notions of difficulty in our setting.

Finally, there is a multi-disciplinary literature on separability (often referred to as decomposition or modularity) of objective functions spanning evolutionary biology, genetic algorithms, and management. This work can be hard to translate into economics, but Page (1996) suggests a measure of difficulty based on decomposing an objective function into as many independent subproblems as possible and reporting the size of the largest subproblem. (For a non-modular objective, this largest subproblem might simply be the original problem itself.) We will argue that under an informal interpretation of this criterion, the bidding problem is approximately separable in 2D but not in the other two treatments.

As we suggested in the introduction, subjects may attempt to decompose the bidding problem into simpler pieces even when this is not a suitable tactic. In a different context, Ashby et al. (1999) confront subjects with a two-dimensional categorization task. They find that subjects tend to rely on simple, one-dimensional classification rules, regardless of whether this is optimal. This is similar to our finding that even though optimal bidding can require tradeoffs across two outcome dimensions (payoffs from winning one unit versus winning two units), our subjects appear to simplify the problem by simply ignoring one of these dimensions.

---

imply that subjects had a better chance to reach equilibrium in our setting; the dynamics within a double auction could promote faster learning, and some of the VBW subjects were experienced. The Durham et al. sessions – 45 to 75 rounds – were longer yet.

3 The Experimental Environment

3.1 Demand and Supplier Costs

There is 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 three suppliers whose roles are played by subjects in the experiment. Each supplier can produce either 0, 1, or 2 units of the good. A supplier's average cost of supplying q units is given by c_q, where

$$c_0 = 0 \qquad c_1 = 100 - \theta \qquad c_2 = \frac{100 + \theta}{2}.$$

A supplier's cost parameter θ, which indexes the convexity of its costs, is private to the supplier. A large value of θ indicates increasing average (and marginal) cost, while a supplier with θ small can supply two units at a lower per unit cost than one unit. If $\theta = 33\frac{1}{3}$, the supplier has constant marginal cost. In the background, one may think of the case $c_2 < c_1$ as a consequence of a large avoidable fixed cost, but this interpretation is not essential to the analysis. In solving for equilibrium, we treat θ as drawn uniformly from $[0, 50]$, but in the experiments we discretize this range to seven evenly spaced types, denoted by $c_1 \in \{50, 58, 67, 75, 83, 92, 100\}$.

This setting is designed to capture, as simply as possible, certain features that make the supply allocation problem interesting. First, there is a mix of efficient production scales – small (large) suppliers minimize per unit cost when producing one (two) units – and none of the cost types is strictly dominated by any other type. Second, an allocation will generally require participation from at least two suppliers at different quantity levels – in this case, two units from one supplier and one from another. Furthermore, an efficient allocation typically requires one large supplier and one small one to win.[6] Partly as a consequence of this, efficiency cannot be determined by at the margin – a cost-minimizing allocation must consider both marginal and inframarginal costs. Using the one-dimensional index θ permits us to incorporate these features in a parsimonious and relatively tractable way.

3.2 Experimental Procedure

The experiments were conducted with undergraduate students in the Netcentricity Behavior Lab at the University of Maryland (UMD) and the vEconlab at the University of Virginia (UVA). There were three treatments, 1U, 2U, and 2D, corresponding to the three auction formats described in the next section. For each treatment, there were five independent sessions (Sessions 1 to 5) with six participants in each session.[7] Details are summarized in Table 1. Each subject earned a show-up fee of $10, which also served as an initial balance to which any profits or losses during the session were added.[8] Earnings during the session were measured in an experimental currency ('francs') and

---

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

[7]Session 5 had more participants – see Table 1 for details.

[8]Subjects were instructed that their final payoff, including this initial balance, would never be negative. While some subjects did make overall losses during the auction rounds, none came close to exhausting his balance. The

Session		1	2	3	4	5
Participants	1U	6	6	6	6	15
	2U	6	6	6	6	12
	2D	6	6	6	6	15
Location		UMD	UVA	UVA	UMD	UMD
Matched?		Yes	Yes	Yes	Yes	No

Table 1: Summary of Experimental Sessions

were later converted to dollars at the rate of 50 francs to \$1 USD. Upon arrival, subjects were given written instructions which were read aloud by the experimenter. Then subjects participated in three practice (non-paying) rounds of auctions to familiarize them with the experimental software.[9] At this point, there was a pause to answer any questions; then the live (paying) rounds began.

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. At the start of Rounds 1, 7, 13, 19, and 25, each subject drew a private cost type independently and uniformly from the seven types shown in Table 2. [10] We refer to cost types $c_1 \in \{50, 58, 92, 100\}$ as *specialized* types and cost types $c_1 \in \{67, 75, 83\}$ as *flexible* types. Informally, we will refer to a supplier as small or large depending on whether $c_1 \leq c_2$ or $c_1 > c_2$, but in certain parts of the data analysis we restrict *small* to be $\{50, 58\}$ and *large* to be $\{92, 100\}$ (with the remainder of types as *flexible*). A subject kept the same cost type for a block of six rounds, called a *sub-session*, before drawing a new type. Both the random matching and the random sequence of cost draws were matched, session by session, across treatments. (That is, for Session i, a full sequence of cost types was drawn and matched into groups of three over thirty rounds. This sequence was then used for Session i of all three treatments.)[11]

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.

smallest final payoff for any subject was approximately \$7.

[9]The experiment was programmed and conducted with the software z-Tree (Fischbacher 2007).

[10]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.

[11]Session 5 is an exception in two respects. First, types and groupings were not matched across treatments. Second, cost types were chosen randomly but not independently in order to ensure balanced representation of different types. 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 independent. This should be taken under consideration (along with other approximations that will be made) when the equilibrium predictions are presented. Non-independence has less bearing on the best response analysis, since this is based on the actual distribution of opponent bids.

Type	1	2	3	4	5	6	7
c_1		50	58	67	75	83	92	100
c_2		75	71	67	63	58	54	50

Table 2: Supplier Cost Types

3.3 Auction Formats

In this section we introduce our three auction formats, the nature of the bidding challenge in each auction and characterize (approximate) equilibrium behavior in these markets. In our characterization of the equilibrium, 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 allocation and competitive bidding.

3.4 1U : One-part bids, uniform pricing (1U)

Under 1U, each bidder simultaneously submits a one-dimensional bid b. 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 uniform price per unit equal to the bid of the marginal winning bidder (that is, the second lowest bidder). For example, if three suppliers submitted bids of 70, 80, and 85, the allocation would be two units for the first supplier and one unit for the second, and both winners would be paid 80 per unit.

Variations on this simple uniform price auction format are common in practice, and when supplier costs are convex, its (theoretical) efficiency can be high, since the uniform price tends toward equating the costs of the marginal unit procured from each bidder. However, with avoidable fixed costs, a supplier's willingness to supply at a particular price may depend on the quantity it will be asked to produce, but there is no way to express this through the bidding. Hence, bidders in 1U confront both *quantity risk* and *price risk*. Quantity risk reflects the fact that any given bid could turn into an obligation to produce either zero, one, or two units, regardless of the intent of the bidder. A large bidder who bids intending to win two units may only win one and fail to cover its costs (illustrating the exposure problem), or a small bidder who intends to be marginal may wind up inframarginal, obligating it to produce an unprofitably large quantity. Bidding lower monotonically improves one's chances of winning two units, and has a non-monotonic effect (first rising, then falling) on the chance of winning one, so there are several different tradeoffs for a bidder to consider as it tries to maximize its expected profit.

A large bidder who succeeds in winning two units always faces some variability in its profits

arising from the fact that its price is set the next lowest rival bid. What is less obvious is that in order to be competitive for two units (and avoid winning one), a large bidder may need to bid below its *two-unit* average cost. Thus, it can end up losing money with its optimal bid, even if it wins its desired quantity, when the price-setting bid is low enough.

This rich set of tradeoffs has two important implications for bidder behavior. First, these incentives interact in such a way to preclude the existence of *any* pure strategy equilibrium – any equilibrium involves mixing by some cost types (discussed below).[12] Second, it allows us to test the implications of risk aversion in a much more nuanced way than is typically possible.

An analytical characterization of the mixed strategy equilibrium is intractible, but we can compute equilibrium strategies numerically for both the continuous and discrete type cases. Figure 1 shows these strategies (presented as cumulative distribution functions over bids) for the discrete case.[13] The smallest suppliers ($c_1 \in \{50, 58\}$) concentrate at $b = 69$. The largest supplier types ($c_1 \in \{92, 100\}$) mix over a compact range of bids. However, each flexible type mixes over two separate intervals, one with low bids and one with high bids.

If one of these flexible types were to always bid in the higher of these two intervals, then market prices would tend to be higher, and any individual supplier of this type would be tempted to try to win two units at a high price by deviating to a low bid. Conversely, if one of these flexible types were to always bid in the lower of the two intervals in its support, then prices would tend to be lower, reducing inframarginal profits. In this case, an individual of this type could do better by abandoning hope of winning two units and deviating to a high bid that might set the price. Thus equilibrium requires these types to mix between bids that are more competitive for the inframarginal and marginal units, respectively.

3.5 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 submits a two-part bid (b_1, b_2) indicating the minimum price *per unit* it is willing to be paid if it supplies one unit (b_1) or two units (b_2). 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. If there is only one subset of bids that achieves this, then it becomes the winning allocation. If more than one subset of bids could provide exactly three units at the market-clearing price, then a tiebreaker is needed. We break ties in favor of the allocation whose cost would be lowest if

---

[12]This is true both with the continuous type space, and with the discretization to seven types.

[13]All of the equilibrium and best response computations in the paper were performed in Matlab; the programs are available on request. In this case, the equilibrium was computed by picking an arbitrary set of starting bid distributions and then computing iterated best responses until the distributions converged. While we cannot prove that this equilibrium is unique, we did try a number of different initial distributions, and all of them converged to these strategies.

the suppliers were paid as bid.[14, 15] By introducing quantity-dependent bidding without quantity-dependent pricing, 2U is in some sense creating two markets but leaving one of them unpriced. This makes bidding less risky, since a supplier can guarantee that its costs will be covered, but it also makes the auction price a less effective signal to suppliers about how to bid successfully.

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.[16] The symmetric equilibrium bids, as a function of one-unit cost c_1 are

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

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

Note that one unit bids are above 90 for all types and all types between $c_1 = 67$ and $c_1 = 100$ pool at 100. Meanwhile, two 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.

---

[14]To illustrate, suppose that the three suppliers submit bids of $(70, 77)$, $(75, 71)$, and $(100, 55)$. At prices below 70, it is only possible to procure a total of two units, both from the third supplier. At a price of 70, it is possible to procure one unit from the first supplier and two units from the third, so this is the winning allocation, and the market-clearing price is 70. Now change the second supplier's bid to $(75, 60)$, holding the other bids fixed. At prices in $[60, 70)$ it is possible to procure either two units (from either the second or third supplier) or four units (two from both), but not exactly three. At a price of 70, there are two ways to procure three units: one unit from the first supplier and two units from either the second or the third. The market-clearing price is again 70, and the tie is broken in favor of the allocation including the third supplier $(70 + 2\,(55) < 70 + 2\,(60))$. This outcome involves passing up a bid below the market-clearing price (the second supplier's two unit bid) in order to procure exactly the quantity demanded by the buyer. Notice that the second supplier could enter the winning allocation here by undercutting the third supplier's two unit bid, and that this would have no effect on the market-clearing price.

[15]This allocation rule – consistent with the instructions provided to subjects – precludes the possibility that the buyer accepts a two unit bid but asks its bidder to deliver only one unit. We thank a referee for pointing out that if the buyer did have this option, then it might sometimes be advantageous to procure three units by accepting two different two unit bids, paying for all four units, but only taking delivery of three. While the equilibrium strategies would certainly change in this case, the chance of being paid for an unproduced and undelivered second unit would still tend to induce bidding below cost.

[16]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.

This equilibrium does not make format 2U look very attractive: equilibrium efficiency should be low because of the pooling and procurement cost should be quite high. However, our goal is not to identify optimal auctions but to examine how successfully suppliers can master a difficult bidding problem. Format 2U delivers such a problem: to do well, a subject must break free from the intuition that a good bid should be in the neighborhood of a markup above her costs.

3.6 2D: Two-part bids, discriminatory pricing (2D)

Under 2D, a bidder submits a two-part bid (b_1, b_2), 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: a supplier will either have b_1 accepted, supply one unit, and be paid b_1, or it will have b_2 accepted, supply two units, and be paid a total of $2b_2$, or it will have no bid accepted and produce zero. The winning allocation is determined by accepting the combination of bids that procures exactly three units at the least total cost to the buyer, subject to the constraint that at most one bid is accepted from each supplier.[17]

An approximation of symmetric equilibrium bidding strategies, as a function of one-unit cost c_1 are,

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

The bidding strategies above 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.

These bidding strategies leverage the fact that under our cost structure, a supplier's one and two unit bids are unlikely to compete with each other. That is, conditional on being close to the margin of winning one unit with its b_1 bid, a supplier is almost surely not close to the margin of winning with its b_2 bid, and *vice versa*. Consequently, 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.[18]

---

[17]For the earlier example with bids of $(70, 77)$, $(75, 71)$, and $(100, 55)$, the allocation remains the same, but now the first supplier is paid 70 per unit, while the third supplier is paid 55 for each of two units. At the conclusion of the auction, all of the bidders see the price vector $(55,55,70)$.

[18]In our model, this is true in part because the ranking of suppliers from lowest to highest one unit cost is always the reverse of the ranking from lowest to highest two unit cost. Thus in the event that a supplier is competitive for one unit, she is generally not simultaneously competitive for two units. In a pilot experiment, we studied format 2D with two bidders rather than three (but no other changes). In this case, a bidder always wins some quantity, so b_1 and b_2 *can* cannibalize each other. We find some modest evidence that subjects underestimate the chance that reducing b_2 will cannibalize one unit profits.

4 Preliminary Results

4.1 Actual vs. Predicted Outcomes

Our first evidence is on aggregate auction outcomes. For each auction in the data, we record the actual outcome on key variables (efficiency, supplier profits, and total production and procurement costs). Then we compute predicted outcomes that would have occurred (for this cost triple), if the subjects had played the strategies derived in Section 3.[19] Then we average these actual and predicted outcome variables across all auctions within a session to form a session mean. If our subjects play mutual best responses to each other, then these actual and predicted session means should be close. We test this by comparing (for each treatment) the five matched pairs of session means, using both a non-parametric Wilcoxon signed rank test and a paired t-test.[20]

The values are in francs per auction; to convert to dollars per session, multiply by 0.6 (30 auctions/session · \$0.02/franc). Note that the statistical tests are conservative – by taking session means as the data points, we do not make any assumptions about the independence of outcomes within a session – but as a result they have relatively low power. In particular, for the Wilcoxon test with $n = 5$, the strongest possible rejection of no difference between actual and predicted outcomes is a p-value of 0.0625; we will refer to this as a significant difference, and to the next smallest p-value (0.125) as marginally significant.

In all three auction formats, subjects earn significantly lower profits than would be predicted by equilibrium; the difference, in dollars per subject per session is \$2.92 for 1U, \$13.54 for 2U, and \$1.43 for 2D. In percentage terms, subjects earn 34% of their equilibrium profits in 2U and 73% in 2D. In 1U, subjects earn -33% of their equilibrium profits since they actually lose money on average. Profits per auction are the residual between the total production cost of the three units sold (SC, for supplier cost) and the total revenue from selling those units (BC, for buyer cost). The shortfall of actual profits can be attributed partly to stronger than expected competition – BC is significantly lower than predicted in all three formats (marginally so for 2D). The other factor to blame for low profits in 1U and 2D is inefficient production – SC is higher for both treatments (significantly for 2D, marginally significant for 1U) than in equilibrium. The last row normalizes suppliers' production cost into an index ranging from 0 (costs were no better than a random allocation) up to 1 (the allocation was efficient).[21] The comparison of actual to equilibrium values of this index is similar to those for SC. Note that in the case of 2U, the fact that subjects

---

[19]For 1U, where the equilibrium strategies for some types are mixed, we compute the expectation of the outcome variable, for each auction.

[20]Note that session means could be close even if *individual* auctions do not conform closely to the predictions. In this sense, a rejection of equality in our tests is a strong refutation of the predictions. However, we will interpret the results cautiously because the predictions are based on strategies that are approximate joint best responses for the discrete cost type case.

[21]More precisely, the index is $Efficiency\,Index = \frac{SC_{rand} - \widetilde{SC}}{SC_{rand} - SC_{eff}}$, where $\widetilde{SC}$ is either the actual supplier cost, or the equilibrium prediction of supplier cost. The expected supplier cost under a random allocation is generated by averaging over the seven possible allocations – six combinations involving two units from one supplier and one from another, and one allocation with one unit from each. We thank a referee for suggesting this measure of efficiency.

do not play the pooling equilibrium strategies means that allocations are more sorted by cost (and so, more efficient) than predicted.

		Treatment		
		1U	2U	2D
Π	Actual	-1.22	11.83	6.53
	Eqm	3.65	34.4	8.91
p-value	Wilcoxon	0.0625	0.0625	0.0625
	(t-test)	0.0067	< 0.001	0.016
SC	Actual	190.06	176.47	172.69
	Eqm	183.04	185.22	170.68
p-value	Wilcoxon	0.125	0.0625	0.0625
	(t-test)	0.037	0.0072	0.0055
BC	Actual	186.41	211.96	192.27
	Eqm	193.99	288.41	197.42
p-value	Wilcoxon	0.0625	0.0625	0.125
	(t-test)	0.0029	< 0.001	0.047
Efficiency Index	Actual	0.36	0.8	0.91
	Eqm	0.56	0.52	0.98
p-value	Wilcoxon	0.125	0.0625	0.0625
	(t-test)	0.042	0.01	0.027

Table 3: Actual Behavior versus Equilibrium Predictions (francs per auction) Π is the (avg.) Profit per supplier, SC the (avg.) total cost of the suppliers selected to produce, BC the (avg.) total cost of the buyer, and Efficiency $= 1$ if the efficient suppliers were selected.

Table 3 suggests that for auction format 2D, equilibrium predicts actual session level outcomes reasonably well. Although there are statistically significant gaps between actual and predicted outcomes, those gaps are small in magnitude. In the other two treatments, the gaps between equilibrium predictions and actual outcomes are substantially wider. It is conceivable that these gaps could narrow in later rounds as subjects gain experience with the auction format. To check for this, we computed the profit shortfall (predicted minus actual profits) for rounds 1-15 and 16-30 of each session, normalized to dollars per 15 rounds. For auction format 2D, the average (across sessions) shortfall does decline over time, from $0.85 to $0.58. (The sum of these is the $1.43 reported earlier.) However, in the other two auction formats there is little change: in 1U the shortfall declines from $1.49 to $1.43, while in 2U it rises from $6.69 to $6.82.[22] Thus the condition

---

[22]Each of the three average shortfalls reported for rounds 16-30 is different from zero with at least marginal

that subjects play mutual best responses appears to fail persistently in some of our settings.

Figure 2 breaks down the actual and equilibrium predicted subject payoffs by cost type.[23] Remember that in our setting, an efficient allocation usually requires the smallest and largest realized suppliers to produce one and two units respectively. The median realized type is what VBW (1996) call a "marginal efficient seller" at both scales of production: he is second-best at producing both quantities that are needed and should sit out. VBW argue that in their setting, these marginal sellers compete too vigorously, lose money on average (due to exposure problems), and contribute to volatile market outcomes. In our setting, it is the flexible cost types who are most likely to be the marginal seller. In 1U (where exposure problems are most severe) these types indeed lose money on average, suggesting that they may contribute to volatile outcomes by bidding too aggressively. However, large bidders also lose money in 1U, and the profit shortfall (of actual profits relative to equilibrium predictions) is shared across all cost types in 1U and 2U. In contrast, in 2D, suppliers with declining average costs (the large and flexible types) earn average payoffs quite close to the predicted ones.

However, Figure 2 cannot tell us whether, for example, large suppliers in 1U appear to do poorly because of their own errors, or because their profit opportunities are limited by other bidders' mistakes. (For example, if small suppliers bid too aggressively, then prices fall and the inframarginal profits that large suppliers rely on dry up.) In the remainder of the paper, we try to separate these issues by focusing on whether subjects play best responses to the empirical opponent bidding that they actually faced.

4.2 Actual vs. Best Response Bidding

In this section, we construct the payoffs that a subject could have earned by best responding to the distribution of her opponents' past play; then we ask how subjects' actual payoffs compare to these computed best response payoffs. This can provide a clearer test of how effectively subjects optimize than the benchmark profits in Table 3 do, since the latter may have been simply impossible to achieve given opponents' actual bidding.

The rationale for computing best responses with respect to *past* play is that it would be both feasible and reasonable for a subject to base her expectations about her current opponents' bidding on her own past market experience. There are two caveats to this approach. One is our implicit assumption that subjects can draw reasonably accurate inferences about past opponent bids from market feedback. Thus success in best responding may partly reflect the richness of market feedback, and subjects' success in interpreting it. The second caveat is that in a volatile market

---

significance (p-value 0.125) in a Wilcoxon test over the five session means. We also tested the difference $Shortfall_{1-15} - Shortfall_{16-30}$ for the five sessions of each treatment to determine whether there is a significant improvement in the gap between actual and predicted profits. There is no significant improvement (for any treatment).

[23]In the figure, actual profit of small bidders in 1U, is calculated by pooling all individual auction payoffs of bidders of type 50 or 58 in all five sessions, and taking the average. Predicted profit is computed, similarly to Table 3, by computing expected payoffs under equilibrium strategies for those same auctions (all realized triples including a small supplier) and averaging. (Other types and treatments are handled similarly.) Statistical comparisons across cost types are deferred until Section 5.1.

environment, a best response to past play might not be a particularly good strategy in the current round. Later in this section, we will show that subjects would have systematically improved their payoffs if they had always chosen best responses to past play. Thus, using these best responses as a benchmark for subjects' behavior seems reasonable.

When an arbitrary bid b is referred to in what follows, it should be interpreted as either a scalar or a pair, depending on the treatment. Define $\pi_{it}(b)$ to be the payoff that i would have earned by playing b under her round t costs, against her opponents' round t bids. Let b_{it} be the bid she actually used, and $\pi_{it} = \pi_{it}(b_{it})$ the payoff she actually earned.

Next we construct the *predicted payoff given past play*, $H_{it}(b)$, defined as the expected payoff that i would earn by playing b in round t, under her round t costs, against the distribution of opponent bids that she faced in rounds 1 to $t-1$. To do this, first define $\pi_{i\tau}^t(b)$ to be the payoff that i would have earned by playing b against her round τ opponents, *under her round t costs*. Then $H_{it}(b)$ is defined as:

$$H_{it}(b) = \frac{1}{t-1} \sum_{\tau=1}^{t-1} \pi_{i\tau}^t(b) \ .$$

We say that a subject best responds to past play if she chooses a bid in round t that maximizes $H_{it}(b)$. We define $\bar{H}_{it}$ to be the payoff that she would actually earn in round t by best responding:[24],[25], [26]

$$\bar{H}_{it} \equiv \pi_{it}(b^*) \quad \text{where } b^* = \arg\max_{b \in B} H_{it}(b)$$

We start by asking how actual profits compare to what subjects could have earned by best responding to past play. We determine how much more money a subject could have earned over 30 rounds by best responding to past play, relative to subjects' actual earnings. The average value of $\bar{H}_{it} - \pi_{it}$ for each treatment, converted to dollars per 30 rounds, is \$1.64 for 1U, \$3.03 for 2U, and \$0.52 for 2D.[27] Thus, subjects leave less money on the table than the comparison to equilibrium payoffs in Table 3 would suggest. For 2D, these foregone profits are a relatively small fraction (14%) of the average 30 round profit of \$3.84. However, the foregone profits represent 42% of average actual profits in 2U (\$7.25), while in 1U the foregone profits would have saved subjects from losing money on average (-\$0.52).[28] Furthermore, these foregone profits persist; over the last 15 rounds they average \$1.65 (1U), \$2.85 (2U), and \$0.42 (2D) respectively (still measured in \$/30 rounds).

---

[24]If $H_{it}(b)$ has more than one maximizer – say b_1^*, b_2^*, ... – then we average over $\pi_{it}(b_1^*)$, $\pi_{it}(b_2^*)$, ... to get $\bar{H}_{it}$.

[25]As a secondary benchmark, we also computed a 'full information' best response. $FI_{it}(b)$ is the expected payoff that b would earn (under round t costs) against the distribution of opponent bids that i faces in all 30 rounds. Similarly, $\overline{FI}_{it}$ is the round t payoff that i would have actually earned by choosing a maximizer of $FI_{it}(b)$. When past and future opponent bidding differ substantially, profits $\overline{FI}_{it}$ and $\bar{H}_{it}$ will tend to differ; to achieve profit $\overline{FI}_{it}$, a subject would need to be good at forecasting opponents' future bids. In our data, $\overline{FI}_{it}$ tends to be higher than $\bar{H}_{it}$, while actual payoffs generally fall short of $\bar{H}_{it}$. Given this, our analysis focuses on the less demanding benchmark.

[26]Of course, there is no guarantee that $\bar{H}_{it}$ will be the largest possible profit for choice it, since it is determined on the basis of opponents' past, not current play.

[27]In each case, we average $\bar{H}_{it} - \pi_{it}$ over all choice situations it, then multiply by $\frac{3}{5}$ = (30 rounds) / (50 francs/\$).

[28]For four out of five sessions in treatments 1U and 2U, the mean value of $\bar{H}_{it} - \pi_{it}$ is positive (and larger than

However, the average foregone profit over an entire treatment does not provide a full picture of how closely actual profits match up with $\bar{H}_{it}$ on a round by round basis. To assess this, we next calculate the mean of the squared deviation $\left(\bar{H}_{it} - \pi_{it}\right)^2$. This value (in francs per choice situation) is 322 for 1U, 540 for 2U, and 61 for 2D.[29]Thus, auction by auction, actual profits π_{it} track $\bar{H}_{it}$ reasonably well in 2D, but less well in 1U and 2U.

Altogether, this preliminary analysis suggests that best responding to opponents' past play is a plausible model of actual behavior in treatment 2D, but does not describe actual behavior in 1U and 2U as well. In the remainder of the paper, we explore these deviations from best response behavior, argue that they can be attributed to imperfections in profit maximization, and assess explanations that could explain these imperfections. We favor explanations that can fit into a general model of bidding across all three treatments, rather than explanations that would require fundamentally different behavior in different treatments. 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 (alternative preferences) would be 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 (hard maximization) is that subjects do care about profits, but their profit maximization problem is challenging, and its "degree of difficulty" varies across auction formats. We find only mixed 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, and hence should be considered when evaluating the tenableness of equilibrium predictions.

5   Bidding Based on Modified Best Responses: Model

In this section, we construct and estimate a model of bidding in the spirit of Erev and Roth's (1998) reinforcement learning and Camerer and Ho's (1999) experience weighted attraction (EWA) learning. Our first objective is to study whether an explanation based on preferences, or one based on noisy optimization, better explains deviations from best response behavior. One criterion of a successful model will be its generality: that is, how well it can explain bidding under all three auction formats.

Suppose that subject i's probability of choosing bid b in round t depends on a propensity $P_{it}(b)$ according to a standard logistic formulation:

$$\Pr\left(b_{it} = b\right) = \frac{e^{P_{it}(b)}}{\sum_{\tilde{b} \in B} e^{P_{it}(\tilde{b})}} \quad , \quad t \in \{2, 3, ..., 30\} \tag{1}$$

---

for the corresponding session of 2D). For both 1U and 2U, Session 1 is the anomaly: $\bar{H}_{it} - \pi_{it}$ is negative (subjects do better than the best response profit, and by a larger margin than for 2D). We will offer a partial explanation for this anomaly later.

[29]In a Wilcoxon test across matched session averages, this mean squared deviation is significantly larger for both 1U and 2U than for 2D ($p = 0.0625$). The difference between 1U and 2U is not significant.

where B is the strategy space. (Subjects' initial bids in round $t = 0$ are left unmodeled.) The propensity is assumed to depend on a linear combination of influences including $H_{it}(b)$ and possibly other factors:

$$P_{it}(b) = \beta H_{it}(b) + [\text{other terms}]. \tag{2}$$

As a baseline, we start with what we will call the *Pure Profit (PP)* model.

$$\text{Model } \mathbf{PP}: \quad P_{it}(b) = \beta H_{it}(b) \tag{3}$$

The previous section presented indirect evidence about whether subjects choose a maximizer of $H_{it}(b)$. Model PP broadens the question to ask how emphatically subjects favor bids with higher profits over bids with lower profits.[30] We begin by testing the hypothesis that a single coefficient β can explain how subjects respond to profits across all three treatments. To do this, we estimate model PP by maximum likelihood separately for each session of each treatment.[31] This generates fifteen (independent) estimates of β. Table 4 presents the five estimates of β for each treatment, and their mean.[32] These five estimates are compared for each pair of treatments with both Wilcoxon signed rank tests and t-tests. While subjects respond more strongly to profits in 2U than in 1U, the difference is not significant. However, subjects do respond to profits much more vigorously in 2D than in the other two treatments, and these differences are highly significant.

Price feedback may partially explain this gap. In treatment 2D, prices reveal all of the winning bids, and this may make it easier for a subject to determine how competitively she must bid in order to win. In the other treatments, the winning inframarginal bid is not revealed by the price; this may contribute to the difficulty that subjects have in identifying profit opportunities.

Next we ask whether the difference between 2D and the other two treatments depends on the supplier's cost type or how experienced she is. We estimate model PP on the pooled data (all 15 sessions) interacting $H_{it}(b)$ with treatment dummies for 1U and 2U (with 2D omitted). We also include interactions with treatment and cost type (flexible or large, with small omitted), or with treatment and round t. The coefficients are presented in Table 5. Small suppliers are equally

---

[30]Note that this model, *via* the definition of $H_{it}(b)$, puts equal weight on a subject's experience in all past rounds 1 through $t - 1$. One might expect that because costs change every six rounds subjects would put greater weight on their most recent experience. (Many learning models, such as Camerer and Ho's EWA learning model, build in a 'decay' factor that can capture this type of effect.) However, when we tested models with a decay factor, we found that the equal-weighting model actually performs better in our data.

[31]For each of these fifteen estimations, each decision opportunity for a subject in rounds 2 through 30 is treated as a single observation. To sidestep computational difficulties related to the size of the strategy space for 2U and 2D, we followed McFadden's procedure (McFadden 1978, Train et al. 1987) of generating consistent estimates by sampling the strategy space. (Further details are in the appendix.)

[32]The standard errors on the coefficient estimates in Table 4 are suppressed. Because we are cautious about treating decisions within a single session as independent, throughout the paper, all standard errors on regression coefficients are robust and clustered at the session level. Sometimes this is impossible (such as here, where each regression includes a single session). In these cases, we try to err on the side of caution by taking session-level outcomes (here, the fifteen estimates of β) as our units of observation for statistical comparisons.

| | Mean | Session | | | | | p-value on test of equality against: | | | |
		1	2	3	4	5		1U	2U	2D
1U	**0.110**	0.175	0.084	0.051	0.132	0.106	signrank	-	0.625	0.0625
							t-test	-	0.392	0.006
2U	**0.160**	0.389	0.213	0.067	0.058	0.075	signrank		-	0.004
							t-test		-	0.0625
2D	**0.391**	0.496	0.557	0.261	0.316	0.325				

Table 4: Session by session coefficients on H_{it} (b) for model PP

H	H*1U	H*2U	H*Flex	H*1U*Flex	H*2U*Flex	H*Lg	H*1U*Lg	H*2U*Lg
0.291	-0.183	0.009	0.025	-0.019	-0.245	0.123	-0.175	-0.348
(0.083)	(0.085)	(0.104)	(0.075)	(0.075)	(0.101)	(0.107)	(0.110)	(0.123)
LL	-15016							

Table 5: Model PP estimated on pooled data, with treatment and cost type interactions (Baseline is treatment 2D). Standard errors are clustered by the 15 sessions.

responsive to profits in 2D and 2U. (In 1U they do worse.) As avoidable fixed costs get larger (the flexible and large interactions), the treatments diverge: profit responsiveness improves in 2D and deteriorates in the other two treatments.

One might wonder whether this performance gap is transient – perhaps subjects just need a few more rounds to sort out best responding in 1U and 2U than they need in 2D. The regression with round interactions refutes this fairly emphatically (Table 6). Responsiveness to profits not only starts at a higher level in 2D, but it also rises far faster with experience in 2D. While the coefficient on H_{it} (b) does rise over time in the other two treatments (at a rate given by the sum of the $H * t$ and $H * t * 1U$ or $H * t * 2U$ coefficients), but the improvement is anemic.

While the results above document clear patterns in how effectively subjects best respond, they do not explain them. In the analysis that follows, we search for parsimonious explanations that can explain behavior in all three treatments well. We focus on two classes of explanation, one based on subject preferences and the other based on identifying auction features that make optimization difficult. Most of the models based on optimization difficulty (payoff volatility, difficult payoff landscapes, salience of dominated strategies) and preferences (risk and loss aversion) are described in Sections 5.1-5.3 and estimated in 6.1-6.3. For methodological reasons, we separate out (in Section 6.5) our study of whether subjects try decompose the profit maximization problem into smaller pieces. Then, in section 7 we investigate the idea that if best responding to past play is difficult, subjects may also rely on simpler heuristics.

20

H	H*1U	H*2U	H*t	H*t*1U	H*t*2U
0.163	-0.091	-0.112	0.016	-0.015	-0.014
(0.031)	(0.036)	(0.033)	(0.003)	(0.003)	(0.003)
LL	-15026				

Table 6: Model PP estimated on treatment and round interactions (Baseline is treatment 2D). Standard errors are clustered by the 15 sessions.

5.1 Payoff Volatility

Earlier, we suggested that avoidable fixed costs and auction rules can interact to produce a volatile environment for bidders; now we introduce measures of that volatility, $V_{it}(b)$, which is a natural counterpart to $H_{it}(b)$. While $H_{it}(b)$ measures the *average* of a sequence of hypothetical payoffs $\left\{\pi_{i1}^t, \pi_{i2}^t, ..., \pi_{i,t-1}^t\right\}$ that b would have earned against past opponents, $V_{it}(b)$ measures the *variance* of that sequence. Formally, define:[33]

$$V_{it}(b) = \frac{1}{t-1} \sum_{\tau=1}^{t-1} \left(\pi_{i\tau}^t(b) - H_{it}(b)\right)^2 \qquad (4)$$

Table 7 presents mean values of $V_{it}(b)$ across the three treatments. Each mean is taken over all possible bids b ($\{0, ..., 100\}$ or $\{0, ..., 100\}^2$) for every choice situation it. The substantially higher values of $V_{it}(b)$ in treatments 1U and 2U (relative to 2D) indicate that in the two uniform price treatments, a typical bid's hypothetical performance in past rounds tends to vary a lot depending on which past round is considered. By averaging $V_{it}(b)$ over all possible bids, the summary statistics in Table 7 include weight on the payoff variability of some truly abysmal bids that subjects are unlikely to seriously consider. In the next line of Table 7, we take the average of $V_{it}(b)$, but this time include only those bids in the top decile of $H_{it}(b)$, for each choice situation it. When we consider the variability of only these 'high payoff' bids, 2U looks more volatile than 1U, but both are still more volatile than 2D.[34]

Based on these summary statistics, it seems plausible that payoff volatility as measured by $V_{it}(b)$ could help to explain deviations from best response behavior in 1U and 2U. We propose two main channels through which $V_{it}(b)$ might affect bidding. The first is based on *preferences*. Suppose that subjects are risk averse and able to estimate both $H_{it}(b)$ and $V_{it}(b)$ reasonably well. When choosing bids, they will tend to penalize riskier ones (high $V_{it}(b)$); then in riskier environments like 1U and 2U, actual bids may diverge relatively more from risk neutral best responses.

Associated with this preference-based argument for deviations from best response behavior,

---

[33]Of course, the statistically correct sample estimate of payoff variance would normalize by $\frac{1}{t-2}$, not $\frac{1}{t-1}$. This would require us to give up the first two rounds of data in the estimation (instead of just the first round). Rather than do this, we use the formula as stated, and accept the implicit restriction that $V_{i2}(b) = 0$ for all b. This has no material effect on the results.

[34]Session by session comparisons of V_{it} in 1U vs. 2D or 2U vs. 2D both reject equality (Wilcoxon p-value 0.0625), regardless of which set of bids V_{it} is averaged over. Differences between 1U and 2U are not significant.

	1U	2U	2D
$\bar{H}_{it} - \pi_{it}$	2.73	5.05	0.86
V_{it} (over all bids b)	438.19	330.18	50.27
V_{it} (over high payoff bids)	158.30	283.02	58.41

Table 7: Summary Statistics for $V_{it}(b)$

we propose the *The Risky Payoff (RP)* model. Under the RP model, we conjecture that subjects dislike risk and bid with the aim of maximizing the simple mean-variance objective function $H_{it}(b) - rV_{it}(b)$, where $r \geq 0$. [35]

$$\text{Model } \mathbf{RP}: \quad P_{it}(b) = \beta \left(H_{it}(b) - rV_{it}(b) \right) \tag{5}$$

The second channel through which $V_{it}(b)$ might affect bidding is based on *noisy optimization.* Define the empirical standard deviation of a bid $\sigma_{it}(b) = \sqrt{V_{it}(b)}$. Suppose that subjects are risk neutral, but their attempts to estimate expected payoffs $H_{it}(b)$ are confounded by the fact that these payoffs are volatile. We suggest two versions of this noisy optimization story; they have similar empirical implications and we will not try to distinguish them in the data. One version is that when a bid's returns have been volatile ($\sigma_{it}(b)$ large), a subject makes larger errors in estimating $H_{it}(b)$. The other is that a subject is able to estimate $H_{it}(b)$ accurately, but like a statistician, places less confidence on this estimate if $\sigma_{it}(b)$ is large. In the first case, a subject may respond strongly to the profit that she perceives, but because this perception is inaccurate, we will measure this as a weaker response to $H_{it}(b)$. In the second case, it is the subject herself who dials down her response because she suspects that the magnitude of $H_{it}(b)$ (either positive or negative) may be large simply by chance.

Associated with this imperfect/noisy optimization-based argument for deviations from best response behavior, we propose the *Noisy Payoff (NP)* model. Under the NP model, we assume that subjects are less sensitive to expected payoffs for bids that are noisy:

$$\text{Model } \mathbf{NP}: \quad P_{it}(b) = (\beta - \gamma\sigma_{it}(b)) H_{it}(b) \tag{6}$$

Note that while both models RP and NP rely on payoff variability, they do so in (testably) different ways: in RP, $V_{it}(b)$ enters directly, whereas in NP, it enters only through the interaction $H_{it}(b)\sigma_{it}(b)$.[36] Furthermore, RP predicts that high payoff variability always makes a bid less attractive, while NP predicts that noise disguises how bad some bids (those with $H_{it}(b) < 0$) are, making them *more* likely to be chosen.

---

[35]In earlier versions of the paper that did not include a learning model, we computed best responses under CARA and CRRA utility and found qualitatively similar patterns across treatments to those that we will report for model RP. Here, with the mean-variance formulation, we err on the side of simplicity.

[36]We use $\sigma_{it}(b)$ in model NP rather than $V_{it}(b)$ partly because of the analogy to confidence interval analysis, and partly because it fits the data better. However, comparisons between models RP and NP do not change if we use $V_{it}(b)$ in both models.

5.2   Bids below cost: loss aversion and weak dominance

As subjects consider the bidding landscape, bids below cost play an interesting role. In two of our auction formats, 1U and 2U, profit maximization may sometimes require placing a bid below cost, which in principle could lose money for some realizations of the opponents' bids (This is not true in 2D, where optimal bids are always above cost.). However, not all bids below cost are potentially profitable. There are also some bids below cost that are dominated and could never be optimal.

There is considerable evidence that people do not treat gains and losses symmetrically (Kahneman et al. (1991)), as manifested in a number of types of behavior which collectively have been termed "loss aversion." If subjects have preferences that penalize loss exposed bids, this could help to explain why profit maximization is apparently weaker in 1U and 2U. It is also possible that some subjects will avoid dominated bids – perhaps because it is particularly salient that they are unprofitable – without avoiding other loss exposed bids.[37] In this case, the behavior suggests not loss averse preferences but rather that dominance helps subjects to optimize more accurately.

Define a bid as loss exposed, with indicator variable $L_{it}(b)$, if there is some set of opponent bids for which it could conceivably return a negative payoff. The values of $L_{it}(b)$ are summarized in Table 8. Table 8 also presents a summary of the bids that subjects actually chose: Note that choosing a loss exposed bid is very common in 1U and rarer in 2U. These summary statistics hint that it may be difficult for loss aversion to simultaneously explain bidding in both 1U and 2U.

Associated with this preference-based argument for deviations from best response behavior, we propose the *The Loss Exposed (L)* model which includes an indicator for loss exposed bids. Under the L model, we conjecture that subjects would avoid submitting bids that are loss exposed if they have loss averse preferences.

$$\text{Model } \mathbf{L}: \quad P_{it}(b) = \beta H_{it}(b) + \gamma_L L_{it}(b) \tag{7}$$

We define another indicator variable $Dom_{it}(b)$ equal to 1 if bid b is weakly dominated. The values of $Dom_{it}(b)$ are summarized in Table 8 in each case, the weakly dominated bids are a subset of the loss exposed ones. Note that choosing a dominated bid is rare in all three treatments, while, as noted above, choosing a loss exposed bid is very common in 1U and rarer in 2U. These combined summary statistics hint that the avoidance of bids below cost may be better explained as a result of (possibly imperfect) profit maximization calculation.

Associated with this imperfect optimization-based argument for deviations from best response behavior, we propose the *The Dominated Bids (Dom)* model which includes an indicator for weakly dominated bids. Under the *Dom* model, we conjecture that subjects would avoid submitting bids that are weakly dominated (but would not avoid bids below cost that are undominated).

$$\text{Model } \mathbf{Dom}: \quad P_{it}(b) = \beta H_{it}(b) + \gamma_{Dom} Dom_{it}(b) \tag{8}$$

---

[37] A reader's first impression may be that the weakly dominated sets for 1U and 2U are obscure rather than self-evident. The underlying logic is less obscure than it may appear. In 1U, a subject should not bid below the marginal cost of her more competitive quantity. (The marginal cost of producing a second unit is $2c_2 - c_1$.) In 2U, the principle is more subtle than "setting both bid components below cost is unwise," but this is a fairly close approximation.

	1U	2U	2D
$L_{it}(b) = 1$ iff.	$b < \max(c_1, c_2)$	$b_1 < c_1$ or $b_2 < c_2$	$b_1 < c_1$ or $b_2 < c_2$
$Dom_{it}(b) = 1$ iff.			
for $c_1 \leq c_2$	$b < c_1$	$\max(b_1, b_2) \leq c_2$	"
for $c_1 > c_2$	$b < 2c_2 - c_1$	$\max(b_1, b_2) \leq c_1$	"
$\frac{\#(\text{chosen bid was loss exposed})}{\#(\text{ choice situations})}$	0.721	0.149	0.015
$\frac{\#(\text{chosen bid was dominated})}{\#(\text{ choice situations})}$	0.105	0.079	0.015

Table 8: Dominated and loss exposed strategies

5.3 Payoff landscape

Another factor that could make optimization difficult for subjects is the shape of the profit landscape. (In this case, there is no alternative based on preferences to compare to directly.) To fix ideas, if $H_{it}(b)$ is strictly quasiconcave in b, we will consider this an easy optimization problem. Concisely characterizing the features of a landscape that make optimization difficult is itself a notoriously difficult problem.[38] We focus on a definition that, while somewhat *ad hoc*, is intuitive and simple. For a choice situation it and a bid b, we say that b is a local maximum if its expected payoff $H_{it}(b)$ is weakly greater than that of every neighboring bid. (For the one dimensional bids in 1U, neighboring bids are $b-1$ and $b+1$. In the other two treatments, we take the four nearest neighbors to $b = (b_1, b_2)$; that is, $(b_1 \pm 1, b_2)$ and $(b_1, b_2 \pm 1)$.) For choice situation it, define LM_{it} to be the fraction of bids (over the strategy space B) that are local maxima.

We conjecture that subjects will appear less sensitive to expected profits in choice situations in which LM_{it} is large. This could occur if subjects search locally for payoff-improving strategies; in this case, they will tend to get stuck at local maxima. Alternatively, subjects may try to take a global view of H_{it} but find it hard to discern patterns in the landscape. Mean values of LM_{it} (see Table 9) indicate that subjects in 1U and 2U tend to face payoff landscapes with many local maxima, suggesting that LM_{it} may help to explain the disparity in best response behavior across treatments.

	1U	2U	2D
$\bar{H}_{it} - \pi_{it}$	2.73	5.05	0.86
LM_{it}	0.29	0.41	0.03

Table 9: Values of LM_{it} across treatments

Associated with this imperfect optimization-based argument for deviations from best response behavior, we propose the *The Local Maximum (LM)* model. As n model NP, we introduce LM_{it} as

---

[38]This question comes up in many fields besides economics (artificial intelligence, genetic algorithms, and numerical analysis, to name a few), but there appears to be no consensus on a simple and robust definition of "difficulty."

an interaction with $H_{it}(b)$:

$$\text{Model } \textbf{LM}: \quad P_{it}(b) = \beta\left(1 - \delta LM_{it}\right) H_{it}(b). \tag{9}$$

In the next section, we report on the success of each of our refined learning models.

6   Bidding Based on Modified Best Responses: Results

One criterion of a successful model will be its generality: that is, how well it can explain bidding under all three auction formats. For each of our paired (altered preferences vs. imperfect optimization) learning models ((RP vs. NP) and (L vs. Dom)) we will compare their performance by two metrics. The first is overall fit to the data, as measured by the log likelihood of the model estimated on the pooled data from all three treatments. This is reported as LL_{pooled} in Table 10. The second metric assesses robustness, or generality, by fitting the model to one treatment and then measuring how well this fitted model predicts choices in the other two treatments. Let $LL_{TonT'}$ be the log likelihood of the treatment T data using the coefficients generated by estimating the model on treatment T'. There are two of these predicted log likelihoods for each treatment T; let LL_{oos}^{T} be their average. (For example, $LL_{oos}^{1U} = \frac{1}{2}\left(LL_{1Uon2U} + LL_{1Uon2D}\right)$.) This measures how well behavior in treatment T is predicted by estimating the model on the other two treatments. To create an overall measure of robust fit, define LL_{oos} to be the sum of these out of sample predictions: $LL_{oos} = LL_{oos}^{1U} + LL_{oos}^{2U} + LL_{oos}^{2D}$. If the model can explain all three treatments well with a single set of coefficients, then LL_{oos} should be close to LL_{pooled}. On the other hand, if LL_{oos} is low, we will argue that the model does not provide a good unifying explanation of bidding behavior.

Here, and with the other models we compare later, our analysis follows a general template. First we compare the models' performance on the aggregate measures LL_{pooled} and LL_{oos}. Then, if one of the models appears less robust, we look at session or treatment level evidence to try to understand why. For this we consult two pieces of disaggregated evidence. One piece of evidence disaggregates LL_{oos}^{1U}, LL_{oos}^{2U}, and LL_{oos}^{2D} session by session. Let $LL_{oos}^{s,T}$ be the contribution of session s to LL_{oos}^{T}, and let $ll_{oos}^{s,T} = \frac{LL_{oos}^{s,T}}{\#(\text{choices } it \text{ in session } s \text{ of treatment } T)}$.[39] This is just the average log likelihood per choice generated for session s of treatment T when the model has been fitted to the other two treatments. The values of $ll_{oos}^{s,T}$ let us see whether behavior in one treatment is predicted particularly poorly by the other two treatments; these values are reported in Figure 4. The second step in evaluating poor robustness is to look at the coefficient estimates when the model is estimated separately for each of the three treatments. This can help to identify the specific cross-treatment difference in behavior that prevents the model from unifying the data well.[40]

---

[39] The reason for the normalization is to be able to compare session 5 (which has more subjects and thus more choice situations) to the other sessions.

[40] To save space, we do not report complete regression tables for these treatment by treatment estimates; the coefficients of interest are reported in the text. The complete tables are available from the authors on request.

	PP	RP	NP	L	Dom	LM	All Pref	All Opt	All
LL_{pooled}	-15644	-15601	-15351	-14554	-14597	-15288	-14550	-14158	-13897
LL_{oos}	-16969	-20377	-16394	-17867	-15999	-17196	-20610	-15742	-22515

Table 10: Log likelihood performance (pooled and out of sample) for the learning models

6.1 Risky Payoffs (RP) vs. Noisy Payoffs (NP)

We start with a comparison of models in which payoff variability affects bidding through either altered preferences RP or imperfect optimization NP (please see Table 10). On both aggregate measures, model NP is more successful than RP. Relative to the Pure Profit (PP) model, NP improves the overall log likelihood by 293 (versus 43 for RP).[41] Furthermore, model NP provides a more robust fit than PP, improving LL_{oos} by 575 log points. The RP model fares much more poorly: adding the preference over $V_{it}(b)$ to the Pure Profit model makes the out of sample likelihood more than 3000 points worse. To investigate the poor robustness of RP, we first consult the disaggregated predicted likelihoods in Figure 4; these suggest that treatments 2U and 2D do not predict risk attitudes in 1U well. (The five circles for LL_{oos}^{s1U} tend to be below the other predictions.) To investigate further, we turn to the treatment by treatment regressions; in these, the raw coefficients on $V_{it}(b)$ are $-r\beta = 0.00068$ for 1U, -0.0059 for 2U, and -0.021 for 2D. In other words, subjects appear to avoid payoff volatility most vigorously in 2D where its level is lowest, and avoid it less vigorously in 2U where volatility is moderate. In 1U, where payoff volatility is high, subjects are actually attracted to riskier bids. Thus RP cannot provide a consistent model of behavior across auction formats; as a result, it performs worse than NP both in overall fit and in robustness.

While the mean-variance risk preferences assumed by model RP are very crude, it is not clear that more sophisticated risk preferences would fit the data better. To explain the tendency to avoid risk more when average payoffs are higher, one would seem to need either increasing absolute risk aversion or context-dependent risk aversion, neither of which is very palatable.

To defend our interpretation of model NP, we take a closer look at bids with negative expected payoffs. Model NP implies that if $H_{it}(b)$ is negative, then b will be chosen relatively more often if it is noisy than if it were not noisy. We test this implication by estimating a version of NP that permits the interaction of σ with payoffs to depend on whether the payoff is positive or negative. The propensity is $P_{it}(b) = \beta H_{it}(b) + \tilde{\gamma}\,\sigma_{it}(b)\cdot H_{it}(b)$, where $\tilde{\gamma} = \gamma_+ \cdot 1\,(H_{it}(b) \geq 0) + \gamma_- \cdot 1\,(H_{it}(b) < 0)$. The estimated coefficients on γ_+ and γ_- are -0.0090 (s.e. 0.0003) and -0.0036 (s.e. 0.00017) respectively. Thus, as predicted by model NP, a bad bid ($H_{it}(b) < 0$) really is more likely to be chosen if its payoff is more volatile.

---

[41]Since model PP is nested in each of the other models in this section, we can compare it to each alternative with a likelihood ratio test. In each case, the improvement in fit under the unrestricted model is highly significant; these results are not reported.

6.2 Loss Exposed Bids (L) vs. Dominated Bids (Dom)

We next compare the models in which subjects assign special significance to some bids below cost either due to loss averse preferences (L) or because these bids are saliently unprofitable (Dom) – see Table 10. In overall fit, both models L and Dom improve fit substantially relative to PP, but L does slightly better (an improvement of 1090 log likelihood points, relative to 1047 for Dom). However, model Dom is substantially more robust; its out of sample likelihood is better than that of model PP, while L does worse than PP.

Figure 4 explains why: model L's predictions explain 2U and 2D well, but 2U and 2D do a poor job of explaining loss attitudes in 1U. Estimating model L separately treatment by treatment confirms this explanation: the coefficients on $L_{it}(b)$ are 0.84 for 1U, -4.21 for 2U, and -5.35 for 2D, so subjects in 2U and 2D appear to avoid loss exposed bids, but subjects in 1U tend to seek them out.[42] In contrast, the coefficients on $Dom_{it}(b)$ in the separate regressions are all negative, indicating that weakly dominated bids are consistently avoided.

The fact that subjects vigorously avoided loss exposed bids in treatments 2U and 2D, but showed much less aversion to them in 1U makes loss exposure unattractive as a robust explanation of behavior; however, the pattern is puzzling and deserves further attention. One possibility is that an indicator for whether a bid is loss exposed does adequately capture loss attitudes. In particular, $L_{it}(b)$ does not reflect whether b is likely or unlikely to lose money. To study this, we estimated the following amended version of model L on the pooled data:

$$P_{it}(b) = \beta H_{it}(b) + \gamma_L L_{it}(b) + \gamma_{\tilde{L}} \tilde{L}_{it}(b) \tag{10}$$

The indicator $\tilde{L}_{it}(b)$ is equal to 1 if $L_{it}(b) = 1$ and $H_{it}(b) < 0$, and equal to zero otherwise. If subjects avoid a loss exposed bid more vigorously when history suggests it is likely to lose money ($H_{it}(b)$ negative), then the coefficient on $\tilde{L}_{it}(b)$ should be negative. In fact, the coefficient is positive and significant ($\gamma_L = -2.66$, s.e. 0.83, $\gamma_{\tilde{L}} = 1.05$, s.e. 0.29). Another possibility is that the pattern has more to do with information processing than loss aversion: confronted with a very large strategy space in 2U and 2D, subjects make the snap judgement that bids below cost can be safely ignored. This serves them well in 2D but is a mistake in 2U. Of course, this explanation is entirely speculative since we have no way to test it in our data.

6.3 Local Maximum (LM)

We finally consider whether the shape of the payoff landscape can help to explain imperfect optimization (Table 10); as noted earlier, there is no alternative based on preferences to compare directly. The LM model improves overall fit a bit more than NP does for the pure profit model

---

[42]This finding for 1U complements certain results in VBW 1996. That paper frames bad market outcomes not as losses, but as small positive payoffs relative to a larger foregone profit from staying out of the market. Thus standard notions of loss aversion have no explanatory power in their setting. Like us, they find that the chance of earning less in the market than could have been earned by sitting out does not discourage subjects from bidding aggressively. The similarity of these two results, despite the fact that we frame bad outcomes as losses and VBW 1996 do not, supports our contention that loss aversion is not a compelling explanation of our subjects' behavior.

(by 356 relative to PP), but its out of sample fit is worse. The improvement in overall fit is not a major surprise. The summary statistics for LM_{it} show that it has some resemblance to an indicator variable for treatments 1U and 2U. We know that the data demand less sensitivity to $H_{it}(b)$ in treatments 1U and 2U, and the interaction term $LM_{it} \cdot H_{it}(b)$ can deliver this. Figure 4 implicates poor predictions from 1U and 2D onto 2U for the weak robustness. Treatment by treatment regressions suggest that treatment 2D is to blame: in 1U and 2U, the interaction term $LM_{it} \cdot H_{it}(b)$ gets a negative and marginally significant coefficients (as expected), but in 2D the coefficient is positive and insignificant. The variable LM_{it} has almost no variation in 2D, so it is no surprise that its effect is imprecisely estimated, but this noise degrades the out of sample predictions based on 2D.

6.4 Comparing Preference-Based and Optimization Based Models

In addition to the models presented above, we also compare a model that combines both preference-based models ($V_{it}(b)$ and $L_{it}(b)$) with one that includes all of the 'difficult optimization' predictors, ($\sigma_{it}(b) \cdot H_{it}(b)$, $Dom_{it}(b)$, and $LM_{it} \cdot H_{it}(b)$); we label these models *All Pref* and *All Ops* respectively. Lastly, we estimate a model, labeled *All* using all six predictors: $H_{it}(b)$, both preference-based predictors, and the three optimization-based predictors (Table 10).[43]

While the *All Opt* model does have one additional free parameter, it substantially outperforms the preference-based specification in overall fit (a log likelihood improvement of 1486 relative to PP, versus 1094). The *All Opt* model is also much more robust across treatments: it improves our measure of robust fit by 1227 relative to PP, while the *All Pref* model does more than 3000 points worse than PP. As Figure 4 indicates, the preference-based model has trouble reconciling 1U with the other two treatments; the way that subjects respond to its predictors in treatments 2U and 2D tends to predict behavior in 1U quite poorly. If one accepts the interpretation that variables $V_{it}(b)$ and $L_{it}(b)$ capture risk and loss attitudes, then these attitudes appear to be highly unstable across auction formats. Unsurprisingly, the model *All* provides the best overall fit to the data, but it inherits the robustness problems of the preference-based models.

6.5 Separability of the Profit Maximization Function

Following the decision-making literature discussed in Section 2, we conjecture that our subjects will be more successful at solving a multi-dimensional profit maximization problem if it can be easily split into smaller pieces. Furthermore, following the results of Ashby et al. (1999), we conjecture that subjects may try to tackle such a problem by reducing it to a simpler one-dimensional problem (even if this is ill-advised).[44]

On the first conjecture, our evidence is very informal. As discussed earlier, in format 2D, a subject does not err much if she treats the auction as composed of independent one-unit and two-unit submarkets. She receives separate price feedback from each of these submarkets, and she can

---

[43]Coefficients on relevant parameters for each model are presented in the appendix.

[44]Because this analysis takes a more disaggregated approach (separating cost types and decomposing $H_{it}(b)$) than the models compared in Table 10, we treat it separately.

tailor one component of her bid to each submarket.[45] In 1U and 2U, this is not possible. In 1U, a supplier is forced to trade off one-unit and two-unit outcomes, since a single bid must deal with both. In 2U, a supplier's one and two unit bids compete with each other more directly due to the uniform price, so the two bid components must be considered as an ensemble. On this basis, we suggest that 2D presents subjects with a more separable optimization problem than the other treatments, and that this helps to explain why subjects appear more responsive to $H_{it}(b)$ in 2D. Obviously this argument is only suggestive; for a rigorous test one would want a more formal definition of separability and an experimental design in which the separability of the bidder's problem varies across treatments independently of other features of the auction.

On the second conjecture, we can be more quantitative. We propose the following two nested hypotheses:

Hyp 1 Suppliers will treat the auction as if it were composed of a one-unit submarket and a two-unit submarket, and focus on profits in one of these submarkets.

Hyp 2 More specifically, a supplier will focus on profits for the quantity at which it is a more competitive producer.

To test these hypotheses, we decompose expected profits into two parts; $H_{it}(b) = H^1_{it}(b) + H^2_{it}(b)$, where

$$H^q_{it}(b) = \frac{1}{t-1} \sum_{\tau=1}^{t-1} \pi^t_{i\tau}(b) 1(\text{b would win q units against i's round } \tau \text{ opponents}).$$

H^q_{it} represents the portion of the expected profit (or loss) for bid b generated by winning q units. We then estimate $P_{it}(b) = \beta_1 H^1_{it}(b) + \beta_2 H^2_{it}(b) + \gamma_{Dom} Dom_{it}(b)$ separately by cost type category (small, flexible, or large) and by treatment.[46] Hyp 1 predicts that in each of these nine regressions, one of the $H^q_{it}(b)$ coefficients will be positive and significant, while the other will be close to zero. Hyp 2 adds the prediction that the coefficient on $H^1_{it}(b)$ ($H^2_{it}(b)$) will be positive for small (large) suppliers. We do not attempt to predict whether flexible suppliers will focus on one or two unit profits. While every flexible supplier has an absolute cost advantage at two units (that is, $c_2 \leq c_1$), her comparative advantage relative to other bidders could go either way. Results are presented in Table 11.

Hyp 1 receives strong support: in eight of the nine regressions, subjects appear to focus exclusively on a bid's potential profits from one of the two quantities, ignoring that bid's potential profits or losses from winning the other quantity. In 1U and 2D, the data also support Hyp 2. Furthermore, in these two treatments, flexible suppliers seem to identify as 'large,' focusing on two-unit profits. However, Hyp 2 is not supported in 2U: all cost types are most responsive to one-unit profits and only weakly responsive to two-unit profits (except, oddly, for small suppliers).

The results for 2U are somewhat puzzling – in particular, the apparent focus of large suppliers

---

[45]We reiterate that this depends on special properties of the cost structure and the number of units demanded that ensure that different components of a supplier's bid are unlikely to be in competition with each other.

[46]Including Dom_{it} helps to absorb out strategies that were never serious contenders for a subject's attention, regardless of which quantity she hopes to win. The results are similar (but a bit noisier, with Dom_{it} omitted.

	1U			2U			2D		
	small	flexible	large	small	flexible	large	small	flexible	large
H^1	0.291	0.025	0.001	0.329	0.326	0.215	0.244	0.023	0.001
	(0.038)	(0.031)	(0.013)	(0.084)	(0.068)	(0.027)	(0.115)	(0.010)	(0.013)
H^2	0.009	0.150	0.270	0.189	0.007	0.010	0.057	0.508	0.270
	(0.015)	(0.089)	(0.045)	(0.032)	(0.024)	(0.009)	(0.017)	(0.145)	(0.045)
Dom	-0.683	-0.710	-6.885	-0.973	-1.849	-4.110	-4.086	-4.443	-6.885
	(0.686)	(0.282)	(0.686)	(0.492)	(0.626)	(0.371)	(0.532)	(0.796)	(0.686)

Table 11: Effect of quantity-specific expected profit on propensity to bid. (Coefficients from nine treatment/cost type regressions, standard errors in parentheses.

on one-unit profits – but we can suggest a potential explanation related to feedback. In 2U, the winning one-unit bid usually becomes public because it sets the market price. On the other hand, a winning two-unit bid tends to be inframarginal, and thus not publicly announced. (This is always true in equilibrium and about 70% of the time in our data.) If superior feedback about winning one-unit bids means that subjects are able to estimate $H^1_{it}(b)$ more accurately than $H^2_{it}(b)$, this could partially account for the puzzle. A second and more mundane factor is that large bidders ($c_1 = 92$ and $c_1 = 100$) have little or no scope to make positive one-unit profits, so for them, maximizing $H^1_{it}(b)$ is almost entirely a matter of avoiding one-unit losses. But this is easily achieved by simply bidding above cost ($b_1 \geq c_1$).

6.6 Summary of Learning Model Results

At the beginning of the paper, we argued that it can be useful to categorize departures from risk neutral profit maximization depending on whether they appear to be the result of subjects' preferences or the result of difficulties in learning to optimize. In our experiments, if we want to attribute these departures to risk and loss attitudes, then we must be prepared to believe that those attitudes are very sensitive to context. Consequently, models that infer risk and loss attitudes from bidder behavior in one context may not provide much guidance about how they will behave in a different environment. On the other hand, factors that tend to make optimization easier (like being able to rule out dominated bids) or harder (like payoff volatility and challenging payoff landscapes) do help to explain bidder behavior, and their effects seem to be fairly robust across contexts. This demonstrates that understanding the difficulty of a bidder's strategic problem is critical to predicting her behavior. More significantly, it shows that it is possible to construct measures of that difficulty that are helpful in predicting behavior across a variety of environments. The models presented so far share a presumption that subjects attempt to think about the expected payoffs to different strategies, including strategies that they have not used before. In the following section, we examine whether simpler heuristics based on reinforcement of bids that were actually used can help to explain behavior.

7 Quantity Reinforcement Heuristics

Models of reinforcement learning usually propose that a subject's propensity to play a particular strategy rises if it has given her good outcomes in the past. Because reinforcement learning only requires a subject to mentally tally her own past outcomes, compared with the counterfactual exercise of estimating how alternative bids would have performed, it could be a particularly appealing heuristic for subjects when markets are complex.

In this section, we study reinforcement of quantities – that is, does a subject tend to return to bids that have won her a particular quantity in the past? The main reason to focus on reinforcement of quantities rather than reinforcement of profits is that a subject's costs change over time. Thus, the profit that b actually earned for her at some point in the past is not necessarily relevant for her today, but the quantity that b earned is still relevant. Second, we have seen that payoffs are particularly volatile in 1U and 2U. This volatility stems from variation in both the price and the quantity that a bid b receives. Focusing on the quantities that b earned is a crude but simple way to eliminate the price variation, giving the subject a more stable measure of the bid's performance.

We also assume that some of the reinforcement of a bid that a subject actually used spills over to nearby bids that she may not have used. This presumes that subjects expect bids near b to perform similarly to the way that b did. The three quantity reinforcement terms described below will be treated as factors that influence the propensity $P_{it}(b)$; they represent progressively more targeted types of behavior.

7.1 Inertia

The least targeted type of behavior that we model presumes that *all* bids that a subject has used in the past are reinforced, regardless of the quantities that they won. This is referred to as inertia, since it implies that a subject will tend to remain in the neighborhood of bids she has used before. Define:

$$I_{it}(b) = \frac{1}{t-1} \sum_{\tau=1}^{t-1} \left(\frac{K}{K + d(b, b_{i\tau})} \right) \quad, \quad t \in \{2, ..., 30\} \tag{11}$$

The term in the summation controls how reinforcement decays for bids at a greater distance $d(b, b_{i\tau})$ from the bid that a subject actually used. The bid that a subject actually used in round τ gets reinforcement one $(d(b_{i\tau}, b_{i\tau}) = 0)$, while the spillover reinforcement for bids further from $b_{i\tau}$ falls to zero as $d(b, b_{i\tau})$ grows large.[47] The measure of distance between two bids is the natural one for 1U and the sum of componentwise distances for 2U and 2D:

$$d\left(b, \tilde{b}\right) = \begin{cases} \left| b - \tilde{b} \right| & \text{for treatment 1U} \\ \left| b_1 - \tilde{b}_1 \right| + \left| b_2 - \tilde{b}_2 \right| & \text{for treatments 2U and 2D} \end{cases} \tag{12}$$

---

[47]Several other specifications for this spillover were investigated, including linear and exponential decay, but this specification provided the best fit to the data.

31

A tendency toward inertia is fairly common in experimental studies (e.g., Haruvy and Popkowski Leszczy (2009) [48]), and we are agnostic about how it should be interpreted. One possibility is that subjects start the experiment with strong priors about which bids will perform well, and so they update these priors relatively slowly as market outcomes accumulate. A slightly different possibility is that subjects expect 'local learning' to pay off. By this, we mean that subjects expect that (i) strategies that are close to each other will have similar payoffs, and (ii) making small adjustments in the direction of higher profits will eventually lead to an optimal strategy. With stable, quasiconcave payoffs this assumption would be justified, but in our setting with avoidable fixed costs, subjects who rely on local learning may never discover profit opportunities that require a more radical change in bidding.

7.2 Win Targeting

Next suppose that subjects respond to market feedback, in a more targeted way, by reinforcing only those bids that win a strictly positive quantity. Of course, win targeting can be at odds with profit maximization, but it has been suggested elsewhere (e.g. Cassidy (1967)) that some subjects may have a 'joy of winning' preference that inclines them toward bids that are more competitive than a pure profit maximizer would choose. Define:

$$W_{it}(b) = \frac{1}{t-1} \sum_{\tau=1}^{t-1} \mathbf{1}(q_{i\tau} > 0) \left(\frac{K}{K + d(b, b_{i\tau})} \right) \quad , \quad t \in \{2, ..., 30\} \tag{13}$$

This term is almost identical to the inertia term; the only difference is that $b_{i\tau}$ and nearby bids are only reinforced if $b_{i\tau}$ earned a quantity $q_{i\tau} > 0$.

7.3 Competitive Quantity Targeting

Alternatively, suppose that a subject expects to be most profitable when she wins the quantity that she can (currently) produce at a lower average cost. Define:

$$R_{it}^q(b) = \frac{1}{t-1} \sum_{\tau=1}^{t-1} \mathbf{1}(q_{i\tau} = q) \left(\frac{K}{K + d(b, b_{i\tau})} \right) \quad , \quad q \in \{0, 1, 2\}, t \in \{2, ..., 30\} \tag{14}$$

$$QT_{it}(b) = \begin{cases} R_{it}^1(b) & \text{if subject } i\text{'s round } t \text{ costs satisfy } c_1 < c_2 \\ 0 & " & c_1 = c_2 \\ R_{it}^2(b) & " & c_2 < c_1 \end{cases} \tag{15}$$

One could think of the subject as keeping separate mental registers for bids that tend to win one unit and bids that tend to win two units and targeting only those bids that tend to win her

---

[48]The authors examine an environment where subjects must choose one of two auctions in which to submit a bid. They find that subjects exhibit a tendency to continue to bid in a particular auction even when switching and bidding in the competing auction would be more profitable.

current low cost quantity.[49] Of the three reinforcement terms, $QT_{it}(b)$ is the only one that captures some element of profit maximization, *via* the emphasis on keeping production costs low.

7.4 Results on Quantity Reinforcement

As a baseline, we use the specification *All Opt* from the previous section, which includes $H_{it}(b)$ and all of the optimization related variables discussed earlier. We add each of the three reinforcement variables one at a time and estimate overall fit using the maximized log likelihood over the pooled data, as earlier; the coefficients for these pooled regressions are reported in Table 12. Then, we calculate a robust measure of fit LL_{oos} as in the previous section.

Individually, each of the reinforcement variables improves fit dramatically relative to the *All Opt* model as measured by both LL_{pooled} and LL_{oos}. Of the three, inertia provides the best overall fit on both measures. We conclude that subjects have a strong tendency to keep bidding in the neighborhood of any bid they have used before. One possible explanation for this is that subjects quickly form strong and accurate beliefs about which bids are profitable, and then tend to stick to those beliefs throughout the experiment. If this explanation were true, then we would expect the strong coefficient on $I_{it}(b)$ to come at the expense of weaker coefficients on $H_{it}(b)$ and $Dom_{it}(b)$. This is not the case: the effects of profits and weak dominance are almost unchanged when the reinforcement variables are introduced. A second problem with interpreting inertia in terms of strong beliefs is that a subject's costs are changing periodically, so many of the bids that contribute to $I_{it}(b)$ will have been placed when the subject had different costs than she does now. It is difficult explain why a subject would consciously decide that those earlier bids are relevant to her current decision. An alternative interpretation is that what we have called inertia is simply identifying individual heterogeneity in the data. Of course, this is really more of a relabeling of the behavior than an explanation. A third possibility, as suggested above, is that when confronted with a large strategy space, subjects explore it locally, by anchoring to familiar bids and experimenting with incremental changes.

Next we ask which type of targeted heuristic, $W_{it}(b)$ or $QT_{it}(b)$ explains behavior better. We estimate one more pair of models including the All Opt variables, inertia, and either win targeting or quantity targeting. The results are mixed. When inertia is included, quantity targeting performs better in overall fit, but its out of sample fit is worse than any combination of inertia and win targeting. We conclude that there is some evidence for both types of behavior.

The coefficients in Table 12 suggest that subjects' reinforcement behavior is more or less orthogonal to the way they respond to expected profits: none of the All Opt coefficients changes much when the reinforcement terms are introduced. A final point worth noting is that there are differences in reinforcement across treatments (Figure 4). In treatment by treatment regressions, the 1U coefficients on inertia and win or quantity targeting are positive and significant, but about half as large as the coefficients for 2U and 2D. (For example, in the model where $I_{it}(b)$ is the only

---

[49]For type $c_1 = c_2 = 67$, $QT_{it}(b)$ should treat quantity 1 and 2 symmetrically, but setting $QT_{it}(b) = 0$ is somewhat arbitrary. We have looked at the alternative definition $QT_{it}(b) = R_{it}^1(b) + R_{it}^2(b)$ in which both quantities are reinforced for this type. (Note this coincides with $W_{it}(b)$.) The results are qualitatively similar.

reinforcement term, the coefficient on inertia is 6.95 for 1U, 15.04 for 2U, and 17.84 for 2D.) One (speculative) explanation for this is that the instability of outcomes in 1U disrupts the inclination that subjects would otherwise have to play similar bids over and over.

	All Opt	I	W	QT	I,W	I,QT	I,W,QT	I,W,QT $\cdot$ Spec
H		0.235	0.247	0.189	0.241	0.219	0.221	0.190
		(0.030)	(0.030)	(0.026)	(0.029)	(0.030)	(0.029)	(0.029)
$\sigma \cdot H$		-0.0038	-0.0036	-.0026	-0.0038	-0.0034	-0.0034	-0.0030
		(0.001)	(0.001)	(0.000)	(0.001)	(0.001)	(0.001)	(0.001)
Dom		-2.78	-3.05	-2.71	-2.86	-2.70	-2.72	-2.59
		(0.563)	(0.487)	(0.401)	(0.526)	(0.522)	(0.52)	(0.480)
$LM \cdot H$		-0.231	-0.267	-0.222	-0.245	-0.231	-0.233	-0.209
		(0.047)	(0.051)	(0.042)	(0.048)	(0.047)	(0.047)	(0.046)
I		12.10			8.38	9.43	8.95	11.95
		(1.41)			(1.47)	(1.16)	(1.38)	(2.47)
W			15.92		5.60		0.947	1.662
			(1.91)		(0.889)		(1.24)	(1.97)
QT			21.98		8.74		8.29	0.025
			(2.97)		(1.70)		(2.07)	(1.31)
$I \cdot Spec$								-4.88
								(2.21)
$W \cdot Spec$								-10.09
								(3.75)
$QT \cdot Spec$								23.40
								(3.39)
n		3306						
LL_{pooled}		-10675	-10910	-11639	-10610	-10494	-10493	-10249
LL_{oos}		-15742	-12610	-12720	-15516	-12408	-13093	-12841

Table 12: Quantity reinforcement models

7.5 Behavior across cost types: Does supplier size affect bidding strategies

In motivating quantity targeting, we argued that it is relatively straightforward for a subject to keep track of which bids tend to win which quantities, and that trying to win her low cost quantity may be a rough proxy for profit maximization. One would expect this targeting argument to be more compelling if a subject is more specialized – that is, if the difference between c_1 and c_2 is large.

To test this conjecture, define an indicator $Spec_{it}$ equal to 1 if subject i is a small or large

supplier in round t. We repeat the estimation of the model with all of the optimization variables and all three reinforcement terms, but this time we also include the interactions $Spec_{it} \cdot I_{it}(b)$, $Spec_{it} \cdot W_{it}(b)$, and $Spec_{it} \cdot QT_{it}(b)$. We predict a positive coefficient for $Spec_{it} \cdot QT_{it}(b)$ and a negative one for $Spec_{it} \cdot W_{it}(b)$; that is, a more specialized supplier will be particularly focused on winning her low cost quantity and avoiding her high cost quantity. This is strongly confirmed in Table 12. The direct effects on $W_{it}(b)$ and $QT_{it}(b)$ become small and insignificant, indicating that flexible cost types do not show evidence of any targeting behavior, and the interactions on $W_{it}(b)$ and $QT_{it}(b)$ are negative and positive, as expected. The interaction on $I_{it}(b)$ (about which we had no prior expectation) is also negative and significant, indicating that flexible types rely more on inertia than specialized types do.

8 Conclusions

Dealing with synergies such as avoidable fixed costs in procurement is an important challenge for auction design. 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 paper, we study an allocation problem that (because of avoidable fixed costs) has certain 'hard' features. Our focus is on how the hard features of the allocation problem resurface as strategic challenges for bidders; our three auction formats package these challenges for bidders in different ways. In our data analysis, we try to shed light on which types of challenge bidders handle well, and which they do not. It should be emphasized that our objective is not to pick a winner from these three auction formats. These formats were not chosen with an eye toward optimality (in the standard mechanism design sense), and discovering the optimal auction for our specialized setting would not be of much general interest in any case. Instead, these formats were chosen because they showcase strategic challenges that we believe bidders will face in many other complex bidding environments. Our hope is that with a better understanding the factors that affect bidder profit maximization in complex environments, we will be able to provide better guidance about when equilibrium predictions are likely to perform badly.

Even though our auction formats are quite different, we find that a few metrics of optimization difficulty can help to explain deviations from best response bidding in all three settings. In contrast, preference-based stories have difficulty explaining all three settings consistently.

Specifically, subjects appear to make more mistakes in assessing the *expected* returns of bids that have higher payoff variability. Subjects are more successful at maximizing single-peaked payoff functions than payoffs with many local maxima. And subjects seem to adopt a "triage" strategy of focusing on one dimension of their bidding problem, even in auction formats where this is ill-advised. Some of these effects appear to be quite robust across different auction formats. However, when we ask whether subjects simply avoid risky bids *per se*, or bids that expose them to losses, the evidence across auction formats is contradictory. If risk and loss attitudes matter to our bidders, then they

must do so in a very context dependent way. Without a theory of that context dependence, we would argue that making predictions about risk and loss aversion in a new and untested auction format is a problematic exercise.

Finally, we also find evidence that, when confronted with a challenging bidding environment, subjects rely in part on simple quantity reinforcement strategies that do not require estimating expected payoffs. Inertia (an attraction to all past bids) is quite strong in all three treatments; this works to subjects' disadvantage in auction formats where their initial instincts about how to bid are wrong. Subjects also show more targeted behavior: a subject with a cost advantage at a particular quantity will favor bids that have won that quantity in the past.

Without convex costs, hard and fast conclusions are difficult to draw, and there are several natural generalizations and extensions that this paper leaves unexplored. For example, it is not clear whether quantity-dependent price signals will be equally effective if the linkages between "submarkets" are stronger. Price signals have a second role to play in guiding market entry and capacity investment decisions that are often discrete and lumpy. We do not compare how well different auctions provide signals for lumpy investment, but it would be interesting to do so. In contrast with our experiment, procurement often involves repeated competition by the same suppliers. While collusion is of course a concern in such settings, when avoidable fixed costs are present, cooperation among a set of suppliers may have a positive role to play in avoiding ruinous miscoordination. This would be worth exploring in future work.

9 Bibliography

Ashby, G., S. Queller, and P. Berretty (1999), "On the dominance of unidimensional rules in unsupervised categorization," *Perception & Psychophysics*, Vol. 61 (6), pp. 1178-1199.

Ausubel, L. and P. Milgrom, (2002), "Ascending Auctions with Package Bidding," Advances in Theoretical Economics, Berkeley Electronic Press, Vol. 1, No. 1, pp. 1019-1019.

Bereby-Meyer, Y. and B. Grosskopf (2008), "Overcoming the winner's curse: an adaptive learning perspective," *Journal of Behavioral Decision Making*, Vol. 21, pp. 15-27.

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

Camerer, C. and T. Ho, (1999), "Experience-weighted Attraction Learning in Normal Form Games," Econometrica, 67, pp. 827–874.

Cassidy, R., (1967), " Auctions and Auctioneering", Univ. of Calif. Press, Berkeley, CA.

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

Cramton, P. (1998), "Ascending Auctions," *European Economic Review*, Vol. 42, pp. 745-756.

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

Cox, J., B. Roberson, and V. Smith (1982), "Theory and Behavior of Single Object Auctions," in Vernon L. Smith (ed.), *Research in Experimental Economics*, Vol. 2, Greenwich: JAI Press, 1982.

Durham, Y. and S. Rassenti and V. Smith and M. Van Boening and N. Wilcox (1996), "Can Core Allocations be Achieved in Avoidable Fixed Cost Environments using Two-Part Pricing Competition?" *Annals of Operation Research*, Vol. 68, pp. 61-88.

Elmaghraby, W. (2007),"Auctions within the e-Sourcing Events", *Production and Operations Management*, Vol. 16, No. 4, pp. 409-422.

Erev, I. and A. Roth (1998), "Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria," American Economic Review, 88, pp. 848–881.

Fischbacher, U. (2007), "z-Tree: Zurich Toolbox for Ready-made Economic Experiments," Experimental Economics 10(2), 171-178.

Goeree, J. and C. Holt (2007), "Hierarchical Package Bidding: A Paper and Pencil Combinatorial Auction," forthcoming in *Games and Economic Behavior*.

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

Haruvy, E. and Peter T. L. Popkowski Leszczyc (2010), "Price Dispersion and Choice in Online Auctions", *Marketing Science*, in press.

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.

Jones, T. and S. Forrest (1995), "Fitness Distance Correlation as a Measure of Problem Difficulty for Genetic Algorithms," *Proceedings of the 6th International Conference on Genetic Algo-*

rithms.

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

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.

Katok, E. and A. Roth (2004), "Auctions of Homogeneous Goods with Increasing Returns: Experimental Comparison of Alternative 'Dutch' Auctions,", *Management Science*, Vol. 50, No. 8, pp. 1044-1063.

Klemperer, P., (2002), "What Really Matters in Auction Design," *Journal of Economic Perspectives*, Vol. 16, No. 1, pp. 169-189.

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.

McFadden, D. (1978), "Modelling the Choice of Residential Location" in A. Karquist et al., eds., Spatial Interaction Theory and Planning Models, Amsterdam: North-Holland.

Page, S. (1996), "Two measures of difficulty," *Economic Theory*, Vol. 8, pp. 321-346.

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.

Rothkopf, M., and A. Pekec and R. Harstad (1998), "Computationally manageable combinatorial auctions," *Management Science*, Vol.44, pp. 1131-1147.

Train, K. (2009), *Discrete Choice Methods with Simulation*, Cambridge University Press.

Train, K., D. McFadden, and M. Ben-Akiva (1987), "The Demand for Local Telephone Service: A Fully Discrete Model of Residential Calling Patterns and Service Choices," *Rand Journal of Economics*, Vol. 18(1), pp. 109-123.

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

Van Boening, M. and N. Wilcox (1996), "Avoidable Cost: Ride a Double Auction Roller Coaster," *American Economic Review*, Vol.86, No.3, pp.461-77.

Yechiam, E. and J. Busemeyer (2008), "Evaluating generalizability and parameter consistency in learning models," *Games and Economic Behavior*, Vol. 63, pp. 370394.

10   Figures

Figure 1: 1U Equilibrium Strategies (Probability density functions of equilibrium mixed strategies, by cost type). The vertical axis is truncated for clarity; types 50 and 58 put full weight on a bid of 69.

Left bar: **Actual** profit of Small, Flexible, or Large cost types, respectively. (Mean over all auction outcomes for each type.)

Right bar: **Predicted** profit of each cost type. (Generated by applying equilibrium bids to the realized triples of cost types in each auction.)

Figure 2: Actual and predicted supplier profits, by cost type

Figure 3: Learning Models out of sample log likelihood per choice ll_{sT}, session by session (1U: blue circles, 2U: green squares, 2D: red triangles)

Figure 4: Alternative Heuristics out of sample log likelihood per choice ll_{sT}, session by session (1U: blue circles, 2U: green squares, 2D: red triangles)

11 Appendix: 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 b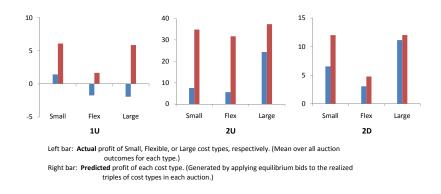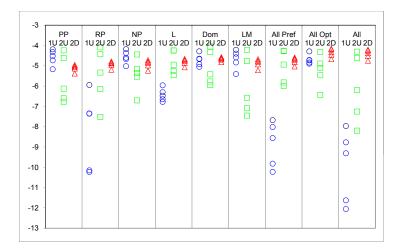idder supplying two units and another supplying one unit). Call this variation 2D'. We claim that when cost types are uniformly distributed on $[0, 50]$, it is an equilibrium of 2D' for each bidder to use the followin g bidding strategies for one and two units respectively:

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

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

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

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

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

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

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

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

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

42

Thus, we must solve

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

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

$$\beta_1^* = \arg\max_{\beta_1} \left(\frac{3}{100}\left(100 - \beta_1\right)\right)^2 \left(\beta_1 - 100 + \theta_1\right) \quad \text{and}$$

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

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 strat egy, 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 $\beta_1 + \beta_2 + \beta_3 = 300 - \frac{2}{3}\left(\theta_1 + \theta_2 + \theta_3\right)$ is smaller than the least cost two bidder allocation. That least cost two bidder allocation involves buying one unit at $b_1^*\left(\theta_{\max}\right)$ and two units at $b_2^*\left(\theta_{\min}\right)$ (where $\theta_{\max} = \max\left(\theta_1, \theta_2, \theta_3\right)$ and $\theta_{\min} = \min\left(\theta_1, \theta_2, \theta_3\right)$), at a total cost of

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

So there is a three bidder allocation whenever

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

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 $[0, 50]$,

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

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}.*

<u>Part 2:</u> 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 $b_1^{*-1}(b_{11}) \lessgtr x$. 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

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

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 $\frac{d}{db}\left(\pi_1(b)\right) = 0$ to get the hazard rate condition:[50]

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

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:

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

---

[50] See below for the second order condition.

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

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

The second derivative is then

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

Notice that $\left(\frac{1 - F(x)}{f(x)} - \frac{2\Delta}{b_1^{*'}(x)}\right)$ 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, $b_1^{*''}(x) = -\frac{100}{(100 - x)^2} < 0$. 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

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

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

Weakly Dominated Bids

1U

$$
Dom_{it}(b) = 1 \text{ iff } \begin{cases} c_1 \le c_2 \text{ and } b < c_1, \text{ or} \\ c_1 > c_2 \text{ and } b < 2c_2 - c_1 \end{cases}
$$

Suppose $\tilde{b} < c_1 \le c_2$. We will show that $\tilde{b}$ is weakly dominated by $\hat{b} = c_1$. Classify opponent bids according to the number of units the subject wins under $\tilde{b}$ (Cases 0, 1, and 2). In Case 0, raising her bid to $\hat{b}$ has no effect on her payoff. In Case 1, her payoff under $\tilde{b}$ is strictly negative; shifting to $\hat{b}$ wins her either 1 unit at a price of c_1 or none, either of which is a strict improvement. In Case 2, write $\tilde{p}$ for the market-clearing price when she bids $\tilde{b}$. If $\tilde{p} > c_1$, then she remains inframarginal under $\hat{b}$, and her payoff doesn't change. If $\tilde{p} < c_1$, then

by shifting to $\hat{b}$ she trades a strictly negative inframarginal payoff of $2(\tilde{p} - c_2)$ for selling one unit at a price of c_1, thus earning zero. (If $\tilde{p} = c_1$, then she trades a weakly negative inframarginal payoff (strictly negative if $c_1 < c_2$) for a lottery between that same negative payoff and a marginal payoff of $\hat{b} - c_1 = 0$.)

Next suppose $\tilde{b} < m_2 < c_1$, where $m_2 \equiv 2c_2 - c_1$. We we show that $\tilde{b}$ is weakly dominated by $\hat{b} = m_2$. Classify opponent bids as above. A switch from $\tilde{b}$ to $\hat{b}$ still has no effect in Case 0. In Case 1, the subject loses money under $\tilde{b}$ (earning $m_2 - c_1 < 0$). The switch to $\hat{b}$ either improves her price on one unit or ejects her from the allocation, both strict improvements. In Case 2, raising her bid to $\hat{b}$ either (i) leaves her still inframarginal, with no effect on her payoff, (ii) shifts her to winning one unit at price $\hat{p} = \hat{b} = m_2$, or (iii) knocks her out of the allocation. In the event of (ii) or (iii), there must be at least one opponent bid no greater than $\hat{b}$, so the original price satisfies $\tilde{p} \leq m_2$. Since $m_2 < c_2$, her original payoff of $2(\tilde{p} - c_2)$ must have been weakly negative, so outcome (iii) is a weak improvement. In the event of (ii), the difference between her new and old payoffs is

$$(m_2 - c_1) - 2(\tilde{p} - c_2) = 2m_2 - 2\tilde{p} \geq 0$$

Thus switching from $\tilde{b}$ to $\hat{b}$ always weakly improves her payoff, as claimed.

2U

$$Dom_{it}(b) = 1 \text{ iff } \max(b_1, b_2) \leq \max(c_1, c_2)$$

We start by showing that a bid satisfying $\max(b_1, b_2) < \max(c_1, c_2)$ is weakly dominated. We will proceed case by case for $c_1 < c_2$, $c_1 = c_2$, and $c_1 > c_2$.

$\underline{c_1 < c_2}$ Suppose $\tilde{b} = \left(\tilde{b}_1, \tilde{b}_2\right)$ satisfies $\tilde{b}_1 \leq c_2$ and $\tilde{b}_2 \leq c_2$. We will show that $\tilde{b}$ is weakly dominated by $\hat{b} = \left(\tilde{b}_1, c_2 + 1\right)$. Classify opponent bids based on whether the subject sells zero, one, or two units under bid $\tilde{b}$ (Cases 0, 1, and 2). Label the market-clearing price under bids $\tilde{b}$ and $\hat{b}$ as $\tilde{p}$ and $\hat{p}$ respectively. Raising a non-winning bid cannot affect the winning price and allocation, so switching from $\tilde{b}$ to $\hat{b}$ has no effect on the subject's payoff in Cases 0 and 1. Consider Case 2. Raising the two-unit bid from $\tilde{b}_2$ to c_2 must weakly increase the market-clearing price, so if our subject continues to sell two units under bid $\hat{b}$, her payoff must weakly rise. Thus we can focus on the case in which raising her two-unit bid shifts her allocation to 0 or 1 unit. Suppose her allocation shifts to 0. This implies that there must be an opponent one unit bid $\leq \tilde{b}_1$ (otherwise, $\tilde{b}_1$ would be accepted in the new allocation). But this opponent one unit bid, together with $\tilde{b}_2 \leq c_2$, imply that $\tilde{p}$ must have been $\leq c_2$, so shifting to quantity 0 must be a weak payoff improvement. Alternatively, suppose that shifting from $\tilde{b}$ to $\hat{b}$ shifts her allocation from 2 units to 1. The fact that her two-unit bid has dropped out of the winning allocation implies that there is at least one opponent's two unit bid (or both of their one-unit bids) that is no greater than $c_2 + 1$. Thus we have $\tilde{p} \leq \hat{p} \leq c_2 + 1$.

The payoffs under $\tilde{b}$ and $\hat{b}$ satisfy

$$
\begin{aligned}
\hat{p} - c_1 - 2(\tilde{p} - c_2) &= 2c_2 - c_1 + \hat{p} - 2\tilde{p} \\
&\geq 2c_2 - c_1 - \hat{p} \\
&\geq c_2 - c_1 - 1 > 0
\end{aligned}
$$

Thus, in this case the subject does strictly better under $\hat{b}$. Thus, $\hat{b}$ weakly dominates $\tilde{b}$.

$\underline{c_1 = c_2}$ In this case, $c_1 = c_2 = 67$. Suppose that $\tilde{b}_1 \leq 67$ and $\tilde{b}_2 \leq 67$, and let $\hat{b} = (100, 67)$. If shifting to $\hat{b}$ does not change the subject's allocation, then her payoff must weakly improve. If her quantity drops from 1 or 2 to 0, then there must be an opponent two unit bid ≤ 67; together with $\tilde{b}_1 \leq 67$, this implies $\tilde{p} \leq 67$. Thus, a quantity drop from 1 or 2 to 0 is weakly payoff improving, as the subject must have been weakly losing money under $\tilde{b}$. Her quantity cannot switch from 2 to 1. ($\hat{b}_1 = 100$ could win only if both opponent one unit bids are 100 and some opponent two unit bid is < 100. But in this case, $\tilde{b}_1$ would win instead of $\tilde{b}_2$.) If her quantity switches from 1 to 2, then she strictly improves: since $\hat{p} \geq \tilde{p}$ and $\hat{p} \geq 67$, $2(\hat{p} - 67) > \tilde{p} - 67$.

$\underline{c_1 > c_2}$ Let $\tilde{b}$ satisfy $\max\left(\tilde{b}_1, \tilde{b}_2\right) \leq \max(c_1, c_2) = c_1$. Write $m_2 = 2c_2 - c_1 < c_2 < c_1$. We consider two cases. If $\tilde{b}$ also satisfies $\max\left(\tilde{b}_1, \tilde{b}_2\right) < m_2$, we will show it is weakly dominated by $\hat{b} = (c_1, m_2)$. Otherwise, it is weakly dominated by $\hat{b} = \left(c_1 + 1, \tilde{b}_2\right)$. (For type $c_1 = 100$, the same arguments go through for $\max\left(\tilde{b}_1, \tilde{b}_2\right) < 100$ using $\hat{b} = \left(100, \tilde{b}_2\right)$.)

First suppose $\max\left(\tilde{b}_1, \tilde{b}_2\right) \geq m_2$. Label $\tilde{p}$, $\hat{p}$, and Cases 0, 1, and 2 as earlier. Since the shift from $\tilde{b}$ to $\hat{b}$ involves increasing the one-unit bid only, it is payoff irrelevant in Cases 0 and 2, so suppose our subject sells one unit under $\tilde{b}$. As earlier, $\hat{p} \geq \tilde{p}$, so if the subject continues to sell 1 unit under $\hat{b}$, her payoff weakly rises. Thus we focus on the possibility that shifting to $\hat{b}$ shifts her allocation to 0 or 2 units. If her allocation shifts to 0, then there must be some opponent $b_2^{opp} \leq \tilde{b}_2 \leq c_1$, that prevents $\tilde{b}_2$ from winning. But this, together with $\tilde{b}_1 \leq c_1$ implies $\tilde{p} \leq c_1$, so she must have been (weakly) losing money under $\tilde{b}$. Thus, dropping to quantity 0 is an improvement. Alternatively, suppose the shift from $\tilde{b}$ to $\hat{b}$ shifts her allocation from 1 to 2. Note that the new price satisfies $\hat{p} \geq m_2$. (The old allocation implies $\tilde{p} \geq \tilde{b}_1$, while the new allocation implies $\hat{p} \geq \tilde{b}_2$. Then $\hat{p} \geq \tilde{p}$ and $\max\left(\tilde{b}_1, \tilde{b}_2\right) \geq m_2$ imply $\hat{p} \geq m_2$.) The difference between the new and old payoffs is

$$
2(\hat{p} - c_2) - (\tilde{p} - c_1) = 2\hat{p} - \tilde{p} - m_2 \geq \hat{p} - m_2 \geq 0
$$

so in this case as well the shift to $\hat{b}$ is a weak improvement over $\tilde{b}$.

Next suppose that $\max\left(\tilde{b}_1, \tilde{b}_2\right) < m_2$ and let $\hat{b} = (c_1, m_2)$, so the shift to $\hat{b}$ involves raising both bid components. This is still payoff irrelevant in Case 0. In Case 1, the arguments are just as in the last paragraph. (The only difference is that if her quantity shifts to 2, $\hat{p} \geq m_2$ is implied directly by $\hat{b}_2 = m_2$.) In Case 2, if she continues to win two units under $\hat{b}$, then her payoff weakly improves. Since her two-unit bid rises, now we must also consider the possibility of a shift to 0 or 1 unit in Case 2. If she shifts to 0 units under $\hat{b}$, then there must be some opponent two-unit

48

bid $b_2^{opp} \leq m_2$. But this implies that $\tilde{p} \leq m_2$, since the combination of $\tilde{b}_1$ and b_2^{opp} would have formed a feasible allocation under $\tilde{b}$. Since $m_2 < c_2$, this implies that our subject must have been losing money under $\tilde{b}$, so earning 0 under $\hat{b}$ is a strict improvement. Finally, suppose toward a contradiction that she shifts to 1 unit under $\hat{b}$ in Case 2. The winning allocation under $\tilde{b}$ would still be feasible now at a price of $\max(\tilde{p}, m_2)$, so we must have $\hat{p} \leq \max(\tilde{p}, m_2)$. Since $\hat{b}_1 = c_1$ is chosen, we also have $\hat{p} \geq c_1(> m_2)$. These imply $\tilde{p} = \hat{p} \geq c_1$. Together this means that the most competitive one and two unit bids by the opponents must be identical, and equal to $\tilde{p}$. Thus, all of our subject's bids – both one and two-unit – must be inframarginal under both $\tilde{b}$ and $\hat{b}$. But the auction rules specify that when a single bidder has two prospective inframarginal bids, the two-unit one is always chosen, so this situation could not arise.

Finally, we must show that no bid satisfying $\max(b_1, b_2) \geq \max(c_1, c_2)$ is weakly dominated. We will be content to show this for the $c_1 < c_2$ case; the logic tranfers easily if $c_1 > c_2$ or $c_1 = c_2$.

Let $c_1 < c_2$, let $\tilde{b} = \left(\tilde{b}_1, \tilde{b}_2\right)$ satisfy $\max\left(\tilde{b}_1, \tilde{b}_2\right) \geq c_2 = \max(c_1, c_2)$, and let $\hat{b} = \left(\hat{b}_1, \hat{b}_2\right)$ be an alternative bid. We will show that there exist opponent bids $\left(b_1^{opp1}, b_2^{opp1}\right)$ and $\left(b_1^{opp2}, b_2^{opp2}\right)$ against which $\tilde{b}$ does strictly better than $\hat{b}$. In the interest of brevity, we omit cases in which $\tilde{b}$ and $\hat{b}$ are not both interior (both bid components belonging to $(0, 100)$). Because there are many cases to consider, the table below tabulates them, in each case indicating the particular opponent bids against which $\tilde{b}$ strictly outperforms $\hat{b}$.

If $\tilde{b}_2 > c_2$ and :

		Opp b_1	Opp b_2		p	q	π
$\hat{b}_2 < \tilde{b}_2$		0	100	Outcome under $\tilde{b}$	$\tilde{b}_2$	2	$2\tilde{b}_2 - 2c_2$
		100	100	Outcome under $\hat{b}$	$\tilde{b}_2$	2	$2\hat{b}_2 - 2c_2$
$\hat{b}_2 > \tilde{b}_2$		0	100	Outcome under $\tilde{b}$	$\tilde{b}_2$	2	$2\tilde{b}_2 - 2c_2$
		100	$\bar{b}_2 = \frac{\tilde{b}_2 + \hat{b}_2}{2}$	Outcome under $\hat{b}$	$\tilde{b}_2$	0	0
$\hat{b}_2 = \tilde{b}_2$	$\hat{b}_1 < \tilde{b}_1,$	100	0	Outcome under $\tilde{b}$	$\bar{b}_1$	0	0
	$\hat{b}_1 < c_1$	$\bar{b}_1 = \frac{\tilde{b}_1 + \hat{b}_1}{2}$	100	Outcome under $\hat{b}$	$\hat{b}_1$	1	$\hat{b}_1 - c_1 < 0$
$\hat{b}_2 = \tilde{b}_2$	$\hat{b}_1 < \tilde{b}_1,$	100	0	Outcome under $\tilde{b}$	$\tilde{b}_1$	1	$\tilde{b}_1 - c_1$
	$\hat{b}_1 \geq c_1$	100	100	Outcome under $\hat{b}$	$\hat{b}_1$	1	$\hat{b}_1 - c_1$
$\hat{b}_2 = \tilde{b}_2$	$\hat{b}_1 > \tilde{b}_1,$	100	$\frac{c_1 + c_2}{2}$	Outcome under $\tilde{b}$	$\frac{c_1 + c_2}{2}$	1	$\frac{c_2 - c_1}{2}$
	$\tilde{b}_1 \leq c_1$	$\bar{b}_1$	100	Outcome under $\hat{b}$	$\bar{b}_1$	0	0
$\hat{b}_2 = \tilde{b}_2$	$\hat{b}_1 > \tilde{b}_1,$	100	0	Outcome under $\tilde{b}$	$\tilde{b}_1$	1	$\tilde{b}_1 - c_1$
	$\tilde{b}_1 > c_1$	$\bar{b}_1$	100	Outcome under $\hat{b}$	$\bar{b}_1$	0	0

If $\tilde{b}_1 > c_2 > c_1$ and :

	Opp b_1	Opp b_2		p	q	π
$\hat{b}_1 < \tilde{b}_1$	100	0	Outcome under $\tilde{b}$	$\tilde{b}_1$	1	$\tilde{b}_1 - c_1$
	100	100	Outcome under $\hat{b}$	$\hat{b}_1$	1	$\hat{b}_1 - c_1$
$\hat{b}_1 > \tilde{b}_1$	100	0	Outcome under $\tilde{b}$	$\tilde{b}_1$	1	$\tilde{b}_1 - c_1$
	$\bar{b}_1$	100	Outcome under $\hat{b}$	$\bar{b}_1$	0	0
$\hat{b}_1 = \tilde{b}_1$ $\hat{b}_2 < \tilde{b}_2$,	100	$\bar{b}_2$	Outcome under $\tilde{b}$	$\bar{b}_2$	0	0
$\hat{b}_2 < c_2$	0	100	Outcome under $\hat{b}$	$\hat{b}_2$	2	$2\hat{b}_2 - 2c_2$
$\hat{b}_1 = \tilde{b}_1$ $\hat{b}_2 < \tilde{b}_2$,	100	100	Outcome under $\tilde{b}$	$\tilde{b}_2$	2	$2\tilde{b}_2 - 2c_2$
$\hat{b}_2 \geq c_2$	0	100	Outcome under $\hat{b}$	$\hat{b}_2$	2	$2\hat{b}_2 - 2c_2$
$\hat{b}_1 = \tilde{b}_1$ $\hat{b}_2 > \tilde{b}_2$,	$\frac{\tilde{b}_1+c_2}{2}$	100	Outcome under $\tilde{b}$	$\frac{\tilde{b}_1+c_2}{2}$	2	$\tilde{b}_1 - c_2$
$\tilde{b}_2 \leq c_2$	100	$\bar{b}_2$	Outcome under $\hat{b}$	$\max\left(\frac{\tilde{b}_1+c_2}{2}, \bar{b}_2\right)$	0	0
$\hat{b}_1 = \tilde{b}_1$ $\hat{b}_2 > \tilde{b}_2$,	100	$\bar{b}_2$	Outcome under $\tilde{b}$	$\tilde{b}_2$	2	$2\tilde{b}_2 - 2c_2$
$\tilde{b}_2 > c_2$	0	100	Outcome under $\hat{b}$	$\bar{b}_2$	0	0

12 Coefficients for Models

Coefficients for the pooled version of each model discussed above are presented in Table 12. These coefficients are all in reduced form. (Thus, for example, a negative coefficient on V or L indicates aversion to risk or loss, while the coefficient on σH identifies $-\gamma$.) Robust standard errors clustered by session are reported, but they should be treated cautiously – they presume no dependence across observations (choice situations it) within a session except for the dependence captured by the model, an assumption that is probably optimistic. With that caveat, all coefficients (with the exception of V in model RP) are significant at the 5% level (and usually at the 1% level as well), and their signs (except for V in the combined preference model) are as expected.

   To get a sense for these coefficients, consider optimization-based models NP, LM and Dom. In model NP, suppose a subject weighs two bids, $\tilde{b}$ and $\hat{b}$ with expected payoff difference $\Delta H = H_{it}\left(\tilde{b}\right) - H_{it}\left(\hat{b}\right) = 10$.[51] If both expected payoffs are noiseless ($\sigma_{it}\left(\tilde{b}\right) = \sigma_{it}\left(\hat{b}\right) = 0$), the model predicts that $\tilde{b}$ is $e^{0.194\Delta H} \approx 7.0$ times more likely to be chosen. Alternatively, suppose both expected payoffs are equally noisy, with $\sigma_{it}\left(\tilde{b}\right) = \sigma_{it}\left(\hat{b}\right) = \sigma$. Then the relative likelihood of choosing $\tilde{b}$ over $\hat{b}$ falls to $e^{(0.194-0.0034\sigma)\Delta H}$. If we plug in the mean value of σ in 1U, 2U, or 2D, the relative likelihood of $\tilde{b}$ over $\hat{b}$ falls to 4.1, 4.6, or 6.1 respectively, reflecting the fact that 1U and 2U are noisier on average than 2D. Suppose that we compare the same two bids, with expected payoff difference $\Delta H = 10$, under model LM. The model predicts that the better of the

---

[51] This is comparable to the average size $|\pi_{it}|$ of a subject's per round profit or loss. The mean values of $|\pi_{it}|$ are 12.6, 13.3, and 6.9 francs in 1U, 2U, and 2D.

	PP	RP	NP	L	Dom	LM	All Pref	All Opt	All
H	0.126	0.125	0.194	0.094	0.066	0.231	0.094	0.237	0.226
	(0.027)	(0.026)	(0.040)	(0.020)	(0.016)	(0.037)	(0.020)	(0.032)	(0.033)
V		-0.0008					0.0002		-0.0001
		(0.001)					(0.000)		(0.0003)
$\sigma \cdot H$			-0.0034					-0.0034	-0.0034
			(0.001)					(0.001)	(0.001)
L				-2.22			-2.27		-1.20
				(0.794)			(0.770)		(0.758)
Dom					-3.38			-3.07	-2.21
					(0.524)			(0.475)	(0.385)
$LM \cdot H$						-0.282		-0.257	-0.228
						(0.064)		(0.049)	(0.046)
n	3306								
LL_{pooled}	-15644	-15601	-15351	-14554	-14597	-15288	-14550	-14158	-13897
LL_{oos}	-16969	-20377	-16394	-17867	-15999	-17196	-20610	-15742	-22515

Table 13: 'Coefficients from the pooled regressions in Section 6. Clustered standard errors in parentheses.

two bids is $e^{(0.231-0.282LM_{it})\Delta H}$ times more likely to be chosen. For a typical choice situation in 2D ($LM_{it} = 0.03$), this relative likelihood is 9.3. For a typical choice situation in 1U or 2U (plugging in the average values $LM_{it} = 0.29$ or 0.41), this relative likelihood falls to 4.4 or 3.2. Finally dominated bids are unpopular: if two bids have equal expected payoffs but one is weakly dominated and the other is not, the weakly dominated one is around 30 times less likely to be chosen. All of these coefficients move in the directions that support an imperfect optimization argument for deviations from best response behavior.

In the preference-based model L, the avoidance of loss exposed bids is similar to avoidance of dominated bids, but somewhat weaker (reflecting the poor performance of $L_{it}(b)$ in treatment 1U). The unstable sign on the direct effect of $V_{it}(b)$ across RP, *All Pref* and *All* provides further evidence that risk aversion does not seem to organize the data well.