Endogenous Reference Price Auctions for a Diverse Set of Commodities: An Experimental Analysis

Olivier Armantier and Charles A. Holt*

December 2015

Abstract

This paper is concerned with multi-object, multi-unit auctions with a budget constrained auctioneer who does not know the value of each object. We propose an allocation mechanism under which the auctioneer i) uses "reference prices" based on value estimates to normalize bids across objects, and ii) reduces value inaccuracies by setting reference prices endogenously using information contained in the bids. We conduct an experiment which confirms that such an endogenous approach increases the auctioneer's profit significantly. These results have important practical implications for many types of auctions, especially in finance.

Keywords: auction design, laboratory experiment, allocation mechanism

JEL: G10, D44, C92, E58

* Federal Reserve Bank of New York and University of Virginia respectively. We wish to thank Paul Klemperer for helpful discussions. We also thank Yunbo Liu, Hannah Murdoch, and Megan Porzio for research assistance. This research was supported in part by the National Science Foundation (NSF 1459918), the University of Virginia Bankard Fund and the Quantitative Collaborative. The views expressed here are those of the authors and do not reflect the Federal Reserve Bank of New York or the Federal Reserve System.

1 INTRODUCTION

This paper is concerned with multi-object, multi-unit auctions with a budget constrained auctioneer. Specifically, the auctioneer receives bids to purchase (or sell) shares of different commodities and must decide which bids to accept conditional on a budget constraint (e.g. a specific amount to raise or to spend). This allocation problem is of great practical importance as many auctions, especially in finance, share these features.

For instance, to control short-term interest rates central banks often drain (or inject) a specific amount of reserves by purchasing (or selling) various types of securities simultaneously at auction.[1] Central banks also routinely conduct liquidity auctions in which a specific amount of funding is allocated against different types of risky collateral.[2] Treasuries occasionally conduct buyback auctions in which a set value of various long-term bonds are purchased simultaneously.[3] Similarly, the 2008 TARP program initially called for the use of auctions to purchase $700 billion worth of different "toxic assets."[4] Financial institutions also commonly raise liquidity by conducting "Bids Wanted in Competition" (BWIC) auctions in which they circulate a list of different loans or bonds for which they are willing to contemplate offers. Finally, to emerge from Chapter 11, a bankrupt company is often asked to put its different assets for sale at auction, with the dual objective of raising liquidity and restructuring the business down to a profitable core (Baird 1993).

While the design of single-object, multi-unit auctions (such as Treasury auctions), and in particular the optimal pricing mechanism (e.g. discriminatory versus uniform price),

---

[1] For instance, the Federal Reserve has most commonly relied on "temporary open market operations" which consist of nearly daily auctions for overnight repurchase agreements in which bids for three classes of securities (Treasury securities, federal agency obligations, and agency mortgage-backed securities) are considered simultaneously for purchase. In 2007, 208 such auctions were conducted for a total purchase size of $1.5 trillion (source: http://www.newyorkfed.org/markets/omo/omo2007.pdf).

[2] Notably, the Bank of England conducts liquidity auctions in which two separate classes of collateral (i.e. more and less risky) are considered simultaneously (see Klemperer 2010, or Frost, Govier and Horn 2015). Liquidity auctions are also conducted by the Reserve Bank of Australia, the Bank of Japan, and the Bank of Canada. Other central banks conducted liquidity auctions on a temporary basis during the great recession. In particular, between December 2007 and March 2010, the Federal Reserve conducted 60 auctions under the TAF (Term Auction Facility) program in which a total of $3.8 trillion in loans were allocated to 429 different financial institutions against collateral with varying degree of risk (see Armantier and Sporn 2013).

[3] In particular, the U.S. Treasury conducted 45 auctions between March 2000 and April 2002 in which $67.5 billion worth of various bonds were bought back (Han, Longstaff and Merril 2007). In each auction, the Treasury set a fixed notional amount to be purchased and could cherry pick among bids for up to 26 different bonds. Other countries have implemented similar buyback programs including Australia, Canada, Finland, Germany, Iceland, Italy, the Netherlands, Norway, Poland, South Africa, South Korea, and Sweden.

[4] These proposed auctions are described and evaluated in Armantier, Holt, and Plott (2013).

has been extensively studied,[5] the literature provides little guidance on how the auctioneer should compare and accept bids across commodities in multi-object, multi-unit auctions. One exception, is Armantier, Holt, and Plott (2013) who proposed a relatively simple approach, a "reference price" auction. A reference price auction is an allocation mechanism, in the class of "scoring auctions," in which the auctioneer first normalizes the bids for each commodity using a value estimate called a "reference price."[6] Then, the highest normalized bids are accepted regardless of the commodity until the auctioneer budget constraint is met. The reference price auction thus enables the auctioneer to accept the bids deemed to have the best relative values across commodities. Of course, the performance of the method is contingent on the quality of the auctioneer's value estimates. The laboratory experiments conducted by Armantier et al. (2013) confirmed that reference price auctions perform well when the auctioneer knows its own value for each commodity, but the auctioneer's profit can be negatively impacted when it has to rely on noisy value estimates.

In this paper, we explore a mechanism to mitigate the impact of noisy value estimates. This mechanism permits reference prices to be adjusted *endogenously* on the basis of the bids submitted. The method rests on the idea that, in general, the auctioneer's values are correlated with bidders' values and, in turn, with submitted bids. Thus, using information contained in the bids should help the auctioneer sharpen value estimates and set more accurate reference prices. An effective mechanism for using the bids submitted to adjust noisy value estimates has to be determined on a case-by-case basis, as it depends on the link

---

[5] See e.g. Back and Zender (1993), Goswami, Noe and Rebello (1996), Nyborg and Sundaresan (1996), Kremer and Nyborg (2004), Abbink, Brandts and Pezanis-Christou (2006), Sade, Schnitzlein and Zender (2006), Brenner, Galai and Sade (2009), or Hortaçsu and McAdams (2010), Morales-Camargo et al. (2013).

[6] Scoring auctions have been used in procurement auctions for quality-based bid adjustments. In the Victoria Bush Tender Auction, for example, a government agency first estimated a "biodiversity index" for proposed land management conservation actions. The bids submitted by farmers were ranked in terms of biodiversity-per-dollar ratios, and the auction was cleared by accepting bids with the highest ratios, subject to the limited procurement budget (Stoneham, et al. 2002). A similar ranking of bids by environmental benefits was used in a laboratory experiment (Cason, Gangadharan, and Duke 2003). The European Space Agency at times solicits bids with price and quality dimensions, which are then ranked according to a score determined by the ratio of quality to bid price (Armantier 1999). Scoring auctions are also used by many states in the U.S. for highway construction contracts. Here, the scored bid is typically determined as a weighted average of the bid price and the road user cost of time delay, which is commonly referred to as "A + B bidding." Asker and Cantillon (2008) present a theoretical analysis of linear scoring auctions and discuss a variety of other applications, including the provision of electricity reserve supply. Chen-Ritzo et al. (2005) use experimental methods to evaluate a reverse auction for procurement in which bid prices are adjusted for lead time and quality. Bidders were told the buyer's marginal values of changes in these non-price dimensions, and were able to adjust their bids during a bid revision period.

between the bidders' signals and the auctioneer's values. We consider a specific scenario in which the bidders have private values drawn symmetrically around the seller's value for each commodity. If competition is sufficient, then the bids should be close to the private values, and the median of the submitted bids should provide a good estimate of the seller's value. Thus, we propose to endogenize each reference price by making it a weighted average of the seller's initial noisy value estimate and the median of the bids submitted for the corresponding commodity.[7]

Given the complexity of the environment (a multi-object, multi-unit scoring auction with budget constraints and endogenous reference prices), it appears difficult to evaluate this endogenous reference price auction directly with economic theory. A commonly used procedure in such situations is to conduct laboratory experiments that simulate the environment with financially motivated human subjects.[8] In this paper, we report on a series of experiments designed to test the extent to which a seller with imperfect value information can benefit from using the bids submitted to set reference prices endogenously.

The environment and the endogenous reference price auction are described in Section 2. We illustrate the approach with a complete information example in Section 3. The experimental design is described in Section 4, and the results are reported in Section 5. A counterfactual analysis is used in Section 6 to evaluate whether alternative procedures for endogenizing reference prices could be more profitable to the seller. Finally, Section 7 reports the outcome of a robustness-check treatment in which subjects are informed of exactly how their bids are used to calculate endogenous reference prices.

2 THE AUCTION ENVIRONMENT

Consider a seller with a portfolio of J commodities. Specifically, the seller owns k_j units of commodity $j = 1, \dots J$. The seller conducts a single auction in order to sell part of its portfolio, subject to a budget constraint that sets sales revenue to be between specified lower

[7] Several alternatives, e.g. taking the *variance* as well as the median of submitted bids into account, are discussed in an exploratory memo written by Goeree and Holt (2008) in consultation with Olivier Armantier.
[8] Experimental methods have long been used to study complex multi-unit auctions, e.g. Cox, Smith and Walker (1984), Goswami, Noe and Rebello (1996), Abbink et al. (2006), Sade et al. (2006), Armantier et al. (2013), or Morales-Camargo et al. (2013). See also Kwasnica and Sherstyuk (2013) for a survey.

and upper limits, Min and Max.[9] The value of a unit of commodity j to the seller is denoted V_j. In some cases, the seller knows V_j before the auction. This is the case for instance in a standard private values auction in which the bidders and the seller receive their own private values for each commodity before the auction. In other cases, the seller does not know its values prior to the auction. This is the case for instance in a standard common value auction in which the value of each commodity is the same for the bidders and the seller, but it is unknown at the time of the auction. In that case, the bidders and the seller receive different noisy signals before the auction. An auction participant can submit multiple bids for each commodity $j = 1, \dots J$, where a bid consists of a price for a unit of a commodity.[10] The seller then compares bids for the J commodities against each other, with the objective of accepting the best offers subject to its budget constraint.[11]

The reference price auction consists of three steps: In step one, the bids for each commodity *j* are transformed into "normalized bids" by dividing them by a reference price R_j reflecting the seller's estimate of the value of commodity j. Observe that a normalized bid above 1 may be considered favorable from the seller's perspective as it exceeds its value estimate. Conversely, a normalized bid below 1 may be considered unfavorable. In step two, the normalized bids for all the commodities are ranked together in a single list from high to low. The bids are accepted in decreasing order of normalized bids, regardless of the commodity, even if the normalized bids are lower than 1, until the lower budget constraint *Min* is reached. In step three, the seller keeps accepting bids in decreasing order of normalized bids as long as i) the bids considered are favorable (i.e. with a normalized bid above 1), and ii) the upper budget constraint *Max* is not reached. The seller stops accepting bids when one of the two conditions is not satisfied.

---

[9] This specification is general as it allows the seller to have a precise liquidity objective, in which case $Min = Max$. In addition, the seller may wish to maintain a balanced portfolio by adding specific budget constraints, Min_j and Max_j, for each commodity j.

[10] Alternatively, we could define a bid as a pair consisting of a price per unit and a number of units.

[11] The environment considered by Klemperer (2010) for his "product-mix" auction is identical, with one important exception: While Klemperer allows the seller's valuation of a commodity to increase with the amount of that commodity sold during the auction, we consider that the seller's valuation remains constant. In this sense, the reference price auction is a special case of the product-mix auction and the two approaches are differentiated only by the assumption made about the seller's preferences. Even when a seller's valuation is not exactly constant, the reference price auction still provides an effective procedure that can easily accommodate a large number of commodities. Note also that the auctioneer knows its valuation functions in Klemperer (2010), so the product-mix auction is not immune to the inaccurate value estimates scenario considered in this paper.

A reference price auction therefore enables the seller to accept the bids deemed to have the best relative values across commodities. The method, however, only characterizes the allocation, not the prices to be paid at the auction. Thus, it needs to be combined with a pricing mechanism such as the standard discriminatory (i.e. pay-as-bid) mechanism or the uniform-price mechanism in which every bidder whose bids are accepted pay the lowest accepted bid (or highest rejected bid) for that commodity. It has now been well established that which pricing mechanism dominates from the auctioneer's perspective is an empirical question that needs to be addressed on a case-by-case basis (Hortaçsu and McAdams 2010).

The experiments conducted by Armantier et al. (2013) suggest that the reference price auction works well when the seller knows its own valuation for each commodity before the auction, in which case $R_j = V_j$. However, as discussed earlier, there are instances in which the seller does not know its own value V_j for commodity j at the time of the auction. In that case, the seller can only rely a noisy value estimate, $\hat{V}_j$. The performance of the auction can be negatively impacted when reference prices are set inaccurately, i.e. when $R_j = \hat{V}_j$. When the value estimate $\hat{V}_j$ is lower than the actual value V_j, the seller is likely to accept unfavorable bids, and conversely, favorable bids are likely to be rejected when the value estimate is too high. In this manner, noisy reference prices reduce the seller's profit (i.e. the difference between the total value sold and the auction revenue).

In this paper, we study a possible method of mitigating the impact of noisy reference prices. To be clear, the objective is not to find an optimal solution that would produce perfectly accurate reference prices. Instead, the objective is to propose an easily implementable method of reducing reference price inaccuracies. The premise is that, in most instances, the seller's and bidders' values should be correlated. Thus, the bids submitted for a commodity should provide the seller with additional information about its own value for that commodity. The seller, therefore, can exploit this information to revise its initial value estimate $\hat{V}_j$. An effective procedure for adjusting reference prices endogenously has to be determined on a case-by-case basis, as it depends on the nature of the link between the bidders' signals and the auctioneer's value.

In the environment considered here, bidders have private values drawn symmetrically around the seller's value for each commodity. If competition is sufficient, then the bids should be close to the private values, and the median of the bids for a commodity should provide a good estimate of the seller's value. We evaluate a procedure that endogenizes each reference price by making it the average of the seller's initial value estimate $\hat{V}_j$ and the median of the bids submitted for the corresponding commodity $\bar{b}_j$: $R_j(\bar{b}_j) = 0.5 * \hat{V}_j + 0.5 * \bar{b}_j$. The median bid is used instead of the mean, since it is likely to be a more robust estimate.

3 AN EXAMPLE UNDER COMPLETE INFORMATION

We now illustrate how endogenous reference prices can benefit the seller when it does not know its own values for each commodity. Consider an auction with 2 commodities, a high quality commodity A and a low quality commodity B. The seller's actual values for A and B are 7 and 3, respectively. The auction has four bidders. Bidders 1 and 2 each want 1 unit of A. Bidders 3 and 4 each want 1 unit of B. As indicated in the first row of Table 1, the values of bidders 1 and 2 for A are 9 and 7, respectively, while the values of bidders 3 and 4 for B are 5 and 4, respectively. The seller has a minimum revenue target of *Min*=7 and a maximum revenue target of *Max*=12. The seller is willing to sell up to 2 units of each commodity using a reference price auction combined with a discriminatory pricing procedure. The seller, however, may not know its own values for the two commodities when setting the reference prices. To simplify, we assume that the seller's values, estimates and reference prices, the bidders' values, as well as the mechanism through which the reference prices are determined are all common knowledge at the time of bidding. Thus, the bidders compete in a complete information environment.[12] Bids are constrained to be integers and ties are decided at random. Finally, to avoid bad equilibria with low bids, the seller imposes a reserve price of 5 for commodity A and 1 for commodity B.

As a benchmark, consider first in the top panel of Table 1 the case of a combined auction without reference prices. In other words, the bids submitted (not the normalized bids) for both commodities are simply ranked together and accepted in decreasing order, subject only to desired revenue targets. In that case, it is easy to verify that bidding the reserve prices is an equilibrium. As indicated in the last three columns of Table 1, the seller's revenue is

[12] Back and Zender (2001) and Sade et al. (2006) also consider complete information to analyze complex multi-unit auctions.

12, the value to the seller of the commodities sold is 20, and the seller's profit -8.[13] This simple auction without reference prices therefore exhibits severe adverse selection in the sense that the lack of sufficient cross-commodity competition enables the bidders to submit low bids at the expense of the seller.

TABLE 1—EXAMPLE UNDER COMPLETE INFORMATION

Seller's Values: 7 for A, 3 for B. Revenue Target: Min=7, Max=12. Reserve Price: 5 for A, 1 for B.

		Security A		Security B		Seller's Revenue	Seller's Value Sold	Seller's Profit
		Bidder 1	Bidder 2	Bidder 3	Bidder 4			
	Bidder's Value	9	7	5	4			
No Reference Prices	Bid	5✓	5✓	1✓	1✓	12 = 5+5+1+1	20 = 7+7+3+3	- 8
	Bidder Profit	4	2	4	3			
Accurate Reference Prices	Bid	8✓	6	4✓	3	12 = 8+4	10 = 7+3	+ 2
	Reference Price	7		3				
	Normalized Bid	1.14=8/7	0.86=6/7	1.33=4/3	1=3/3			
	Bidder Profit	1	0	1	0			
Inaccurate Reference Prices	Bid	6✓	6✓	4	3	12 = 6+6	14 = 7+7	- 2
	Reference Price	4		3				
	Normalized Bid	1.5=6/4	1.5=6/4	1.33=4/3	1=3/3			
	Bidder Profit	3	1	0	0			
Endogenous Reference Prices	Bid	7✓	6	4✓	3	11 = 7+4	10 = 7+3	+ 1
	Value Estimate	4		3				
	Median Bid	6.5=(7+6)/2		3.5=(4+3)/2				
	Reference Price	5.25=(4+6.5)/2		3.25=(3+3.5)/2				
	Normalized Bid	1.33=7/5.25	1.14=5/5.25	1.23=4/3.25	0.92=3/3.25			
	Bidder Profit	2	0	1	0			

✓ indicates a bid accepted by the auctioneer.

Consider now the case of an accurate reference prices auction in the second panel from the top in Table 1. In that case, the seller homogenizes the bids submitted using reference prices equal to its own values. It is easy to verify that bidding the reserve price is no longer an equilibrium (the seller would not accept all bids, giving each bidder an incentive to raise its bid). Instead, by bidding 8 and 4, bidders 1 and 3 are guaranteed to purchase a unit of each commodity and thus have no incentives to submit higher bids. Further, given the bid of bidder 4, submitting lower bids would not increase the profit of bidders 1 and 3. Thus, the bids in the second panel from the top in Table 1 are in equilibrium when the seller has accurate reference prices. As indicated in the last three columns of Table 1, the seller's profit is now positive (+2). This outcome is consistent with the experiment results reported in

[13] Alternatively, the seller could decide to accept only the two highest bids of 5 (as they meet the seller's minimum sales target), which would still result in a seller's profit of -4.

Armantier et al. (2013), in which scoring the bids with accurate reference prices mitigate the problem of adverse selection.

Next, consider the case of inaccurate reference prices in Table 1 (third panel from the top). In this case, the seller does not know its values for commodities A and B. Instead, let us assume that the seller receives value estimates of 4 and 3 for A and B, respectively. Thus, the seller's value estimate is accurate for the low quality commodity B, but undervalued for the high quality commodity A. The seller then sets the reference price for each commodity equal to its value estimate. As indicated in the third panel from the top in Table 1, in equilibrium, bidders 1 and 2 can take advantage of the undervalued reference price on commodity A by submitting low bids of 6, which nevertheless guarantees that they clinch one unit of A each (a lower bid of 5, normalized to 1.25, would lose to the bid of 4, normalized to 1.33, from bidder 3). As indicated in the last column of Table 1, the seller profit is now -2. Thus, the seller allocates two units of the commodity it undervalued at a price below its actual value. This example illustrates how performance might be degraded for the seller when reference prices are noisy. Note, however, that the two bids of 6 for commodity A should have provided the seller a signal that its value estimate of 4 was too low. This is the core idea behind the endogenous reference price presented next.

In the endogenous reference price auction presented in the bottom panel of Table 1, the seller has the same inaccurate value estimates of 4 and 3 for A and B as before. The seller now sets the endogenous reference price for a commodity midway between its initial value estimate and the median of the bids submitted for that commodity. As a result, bidders 1 and 2 are not guaranteed anymore to purchase a unit of A with a bid of 6. Indeed, their median bid (6) would raise the reference price from 4 (the seller's value estimate) to 5 when the reference price is endogenized. This is still below the seller's value of 7, but sufficient to give bidder 1 an incentive to raise its bid to 7. In equilibrium, bidders 1 and 3 are allocated a unit of commodities A and B with bids of 7 and 4 respectively, which results in a seller's profit of +1. The seller's profit is thus lower than when the reference prices are accurate (+2), but substantially higher than when the reference prices are inaccurate and unadjusted (-2).

It is easy to verify that the equilibrium bids in Table 1 remain unchanged if, instead of a discriminatory pricing rule, the seller uses a uniform pricing rule under which the price a winning bidder pays for a commodity is equal to the lowest accepted bid for that commodity. Thus, the seller's profit is exactly the same whether the seller implements a uniform or a discriminatory auction. This result illustrates that the allocation mechanism, not the pricing mechanism, can be of first order importance to maximize the seller's profit.

To summarize, the example illustrates how noisy value estimates can lead to problems that might be partially corrected with endogenous adjustments to reference prices, a possibility that we investigate next by running experiments.

4 THE MAIN EXPERIMENT

The basic setup involves a forward, discriminatory price auction sale with affiliated private values.[14] In each auction, four subjects can bid on four commodities, labeled *A*, *B*, *C*, and *D*. The value of a unit of a commodity to the seller is set to be a "mean market value" that is drawn independently from a uniform distribution on the range [$20, $80]. The affiliated value of a unit of a commodity to each bidder is then drawn independently from a uniform distribution with range of +– $10 around the mean market value.

A bidder can submit up to 6 different bids on each commodity, where each bid consists of a price for a unit of that commodity. Bids are restricted to be integer-valued between a minimum (reserve price) of $10 and the bidder's value. The minimum and maximum revenue targets for the seller are specified to be *Min*=$1,200 and *Max*=$2,000 respectively. These budget constraints were chosen so as to make the auction competitive. Indeed, a rough back-of-envelope calculation suggests that, on average, the auctioneer should accept no more than half the units demanded by bidders.[15]

The allocation and the pricing mechanism, the seller's revenue targets, the distribution of the mean market values, the conditional distribution of the bidder's values, and the number of bidders are all common knowledge. Unlike the example in Section 3, however, the experiment is conducted under incomplete information. In particular, bidders in each auction are informed of their own values for each commodity, but they do not know

[14] In contrast, the experimental setup used in Armantier et al. (2013) to study the basic reference price auction was based on a reverse (procurement) auction, a common value setting, and a uniform price mechanism.

[15] There are 4 bidders per commodity, each demanding 6 units of a commodity, and each unit is worth $50 on average. Thus, the total value demanded per commodity is $1,200 on average and the total value demanded for all 4 commodities is $4,800, which is more than double the maximum revenue constraint of $2,000. Even if bidders bid below value, the auctioneer should not accept more than half of the value demanded.

the other bidders' values. Further, while the subjects are told that the seller would use a reference price auction, the reference prices actually used to score the bids in each auction are not announced prior to bidding. The subjects are only told that the seller would use the "best available information" to determine its value estimates, although the details, e.g. whether the seller's value estimates are noisy, accurate or endogenous, are not disclosed.[16] Thus, in every treatment of the main experiment discussed in this section, the subjects do not know how the reference prices are determined by the seller.

After each auction, the reference prices, normalized bids, and acceptance status for all bids are announced, so that subjects can verify how the allocation is determined. The seller's values (i.e. the mean market values) and the other bidders' values are not revealed, either before or after the auction. Thus, in every treatment of the main experiment, the subjects do not know whether the seller had accurate or noisy value estimates.

The main experiment consists of three treatments. In the "Known" treatment, which serves as a baseline, the reference price for each commodity is set to the seller's value. In the "Noisy" treatment, the seller does not know the mean market values. Instead, for each commodity, the seller receives a noisy signal drawn from a distribution that is uniform on a range of +–$15 around the mean market value. The reference price is set to be this noisy estimate. In the "Endogenous" treatment, the seller receives the same noisy signals as in the Noisy treatment, but the endogenous reference price for a commodity is set to the average of the noisy signal and the median of the bids submitted for that commodity. Note that the instructions given to subjects are the same in all three treatments, so that subjects are unaware of any treatment difference. Each treatment consists of 6 independent sessions, each with 4 bidders and a sequence of 12 auctions with different random draws. The same random draw series is used for all auction sequences, regardless of session and treatment. Thus, each of the 12 auctions is directly comparable across sessions and treatments.

The main experiment involved 72 subjects (6 auction sequences for 3 treatments, with 4 bidders each). Another 24 subjects participated in a robustness check to be described in

[16] The effects of revealing information about endogenous adjustments to reference prices are investigated in a robustness check treatment reported in Section 7. Further, note that in real world multi-object auctions it is common for the auctioneer not to make the allocation process entirely explicit. For instance, in the open market operations and buyback auctions discussed in footnotes 1 and 3, the auction rules do not specify how the auctioneer would select among bids for different commodities.

Section 7. Subjects had not participated previously in reference price auctions. Pilot experiments and software tests were done at the University of Virginia, and the experiment was conducted at the University of Montreal. The web-based Veconlab interface displayed instructions (included in an appendix) and bidding pages. Each session lasted for about an hour. Subjects were paid a show-up fee of $10 (Canadian), with subsequent earnings determined by a fractional payout (0.08) of the subject's earnings for all 12 auctions. Earnings averaged about $35 (Canadian), plus the show-up fee.

5 EXPERIMENT RESULTS

5.1 SELLER'S PROFIT COMPARISON

First, we compare the three treatments with respect to the seller's profit, i.e. the difference between the total sales revenue and the total value to the seller of the units sold at the auction. Summary statistics, aggregated over the 6 independent sessions, are provided for each treatment in the first column of Table 2 (ignore for the moment the last row of Table 2 which will be discussed in Section 7). In addition, the seller's profits (averaged across the 6 independent sessions) for the 12 auctions conducted in each treatment are plotted in Figure 1. Recall that the same random draws are used across all treatments in a given auction, so auction-by-auction comparisons are meaningful.

TABLE 2—SUMMARY PERFORMANCE MEASURES BY TREATMENT

Treatment	(1) Seller's Profit	(2) Reference Price Inaccuracy	(3) Bid Discount
Known	$728.55	0%	11.95%
Noisy	-$569.81	17.6%	13.27%
Endogenous	$194.58	5.8%	11.60%
Endogenous Explained	$564.91	5.3%	9.62%

Notes: In the "Known" treatment (first row) the seller knows its own values and reference prices are accurate. In the "Noisy" treatment (second row) the seller does not know its own values and must rely on noisy estimates to set the reference prices. In the "Endogenous" treatment (third row), the seller has the same noisy estimates as in the Noisy treatment and sets the endogenous reference price as the average of the noisy estimate and the median of the bids submitted. The "Endogenous Explained" treatment (fourth row) is identical to the Endogenous treatment except that the subjects were told how their bids would be used to endogenize reference prices. The "Seller's Profit" is the difference between sales revenue and market value of sales to the seller. The "Reference Price Inaccuracy" is the average absolute percentage deviation between the reference price and the seller's actual value. The "Bid Discount" is the average of the differences between a bidder's value and the bid submitted, expressed as a percentage of value.

Figure 1: Seller's Profit

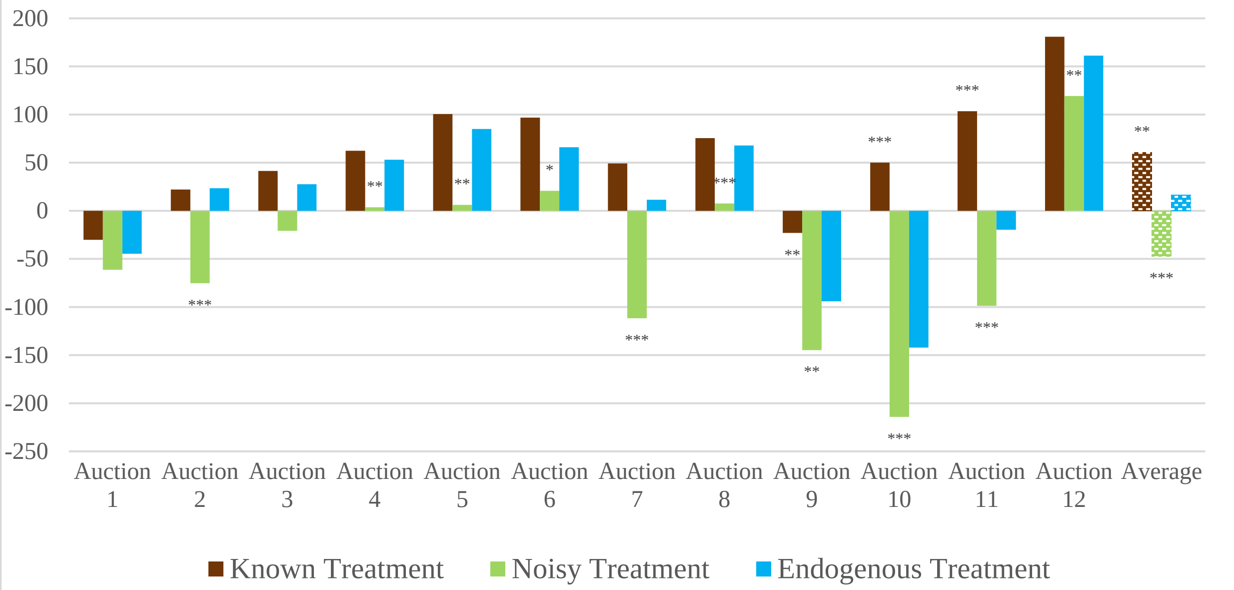

Notes: Each bar in the chart represents the average of the 6 sessions conducted for the corresponding treatment. The same random draws are used in each session and treatment, so the bars in each auction are directly comparable across treatments. The stars above a bar represents the outcome of a 2-tail permutation test, by auction, of equal means between the Endogenous treatment (last bar) and the treatment corresponding to the bar under consideration. The asterisk superscripts, ***, **, and * indicate that the null hypothesis of equal means is rejected at the 1%, 5%, and 10% significance levels.

As indicated in column 1 of Table 2, the seller's profit is, on average, negative for the Noisy treatment. In contrast, the seller's profit in the Known treatment (in which the seller knows its own values) is positive on average and significantly higher than in the Noisy ($p < 0.01$) and Endogenous ($p < 0.05$) treatments.[17] As shown in Figure 1, the seller's profit in the Known treatment is almost always positive and it is always higher than in the other two treatments with Noisy reference prices. This leads to our first result.

> ***Result 1****: The seller's profit is higher when the seller knows exactly how much it values each commodity.*

Observe that *Result 1* is consistent with the results reported in Armantier et al. (2013) for a different auction environment (i.e. a common value, uniform price, reverse auction) and with the intuition presented with the complete information example in Section 3: The performance of the reference price auction is impacted negatively when the seller does know its own values and has to rely on noisy estimates.

[17] Unless otherwise noted, the significance levels are determined by running 2-tail permutation tests on the relevant session averages that are reported in Appendix A.

We now compare the two treatments with noisy reference prices with respect to seller's profit. As indicated in column 1 of Table 2, the Endogenous treatment dominates the Noisy treatment: on average, the seller's profit is negative in the Noisy treatment, but it turns positive in the Endogenous treatment. Further, the seller's profit is significantly higher ($p <$ 0.01) when noisy reference prices are endogenized. Figure 1 confirms that this result is systematic: on average, the seller's profit is always higher in the Endogenous treatment than in the Noisy treatment, and it is significantly higher in 10 of the 12 auctions. Note also that the difference in profit is substantial. Indeed, when compared with the seller's profit in the Known treatment, the total losses caused by noisy value estimates are reduced from $1,298.4 in the Noisy treatment ($728.6 – (–$569.8)) to $534.0 in the Endogenous treatment ($728.6 – $194.6). In other words, the losses due to noisy value estimates are reduced by nearly 60% when reference prices are endogenized. This leads to our second result.

> ***Result 2****: When the seller has noisy value estimates, endogenizing the reference prices by taking the average of the noisy value estimate and the median of the bids submitted increases the seller's profit.*

In the next subsections, we explore the possible factors that may explain why the Endogenous treatment dominates the Noisy treatment with respect to seller profits.

5.2 DOES ENDOGENIZING REFERENCE PRICES REDUCE INACCURACIES?

We explore whether endogenizing the reference prices (by taking the average of the noisy estimate and the median bid) produces more accurate reference prices. The reference prices used to score bids in the Noisy and Endogenous treatments (vertical axis) are plotted in Figure 2 as a function of the seller's actual values (horizontal axis).

Figure 2 indicates that the endogenous reference prices (hollow dots) are generally closer to the seller's values (i.e. the diagonal) than the noisy reference prices (solid dots). Thus, the reference prices used to normalize bids in the Endogenous treatment are more accurate than the reference prices used in the Noisy treatment. This result is confirmed in the second column of Table 2 where we can see that the average absolute percentage deviation between the seller's actual values and the reference prices in the Endogenous treatment (5.8%) is considerably smaller than in the Noisy treatment (17.6%). Further, this difference is significant, based on a 2-tail permutation test ($p < 0.01$). This leads to our third result.

***Result 3**: Endogenizing the reference prices by taking the average of the noisy value estimate and the median of the bids submitted reduces reference price inaccuracies by more than half on average.*

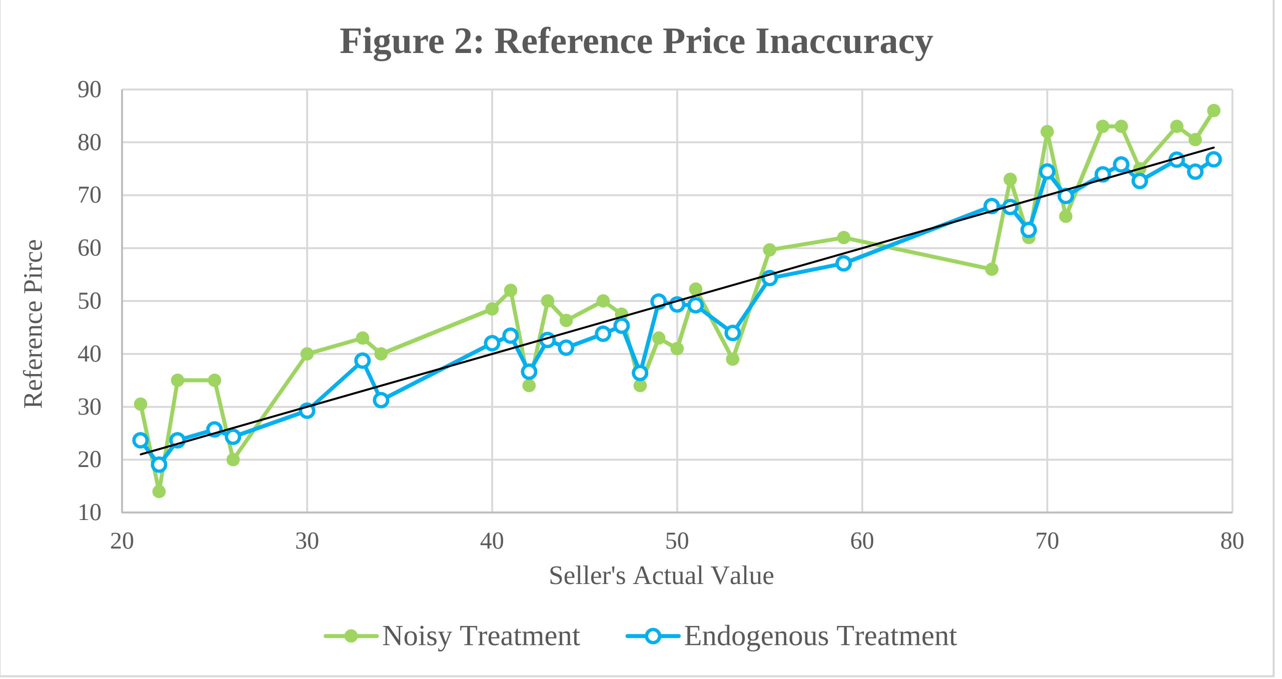

Notes: Reference prices averages are computed by considering all auctions for each realized seller value. The diagonal corresponds to accurate reference prices. Endogenous reference prices are indicated by hollow dots.

5.3 Acceptance Probability

Recall that noisy reference prices should lead the seller to incorrectly accept bids on commodities it undervalues and to incorrectly reject bids on commodities it overvalues. Next we determine whether the reduction in reference price inaccuracies identified in the previous section provides some protection against "incorrect" accept and reject decisions. Figure 3 plots the change in the probability of accepting a bid for a unit of a commodity (regardless of the amount bid) as compared with the Known treatment (vertical axis), as a function of the inaccuracy of the reference price used to score the bids for that commodity in the Noisy and in the Endogenous treatments (horizontal axis). For instance, a horizontal coordinate of 15% and a vertical coordinate of –40% indicates that the seller is 40% less likely to accept a bid (compared with the Known treatment) on a commodity whose reference price has been overvalued by 15%.

Figure 3: Change in the Probability of Accepting a Bid as a Function of the Reference Price Inaccuracies

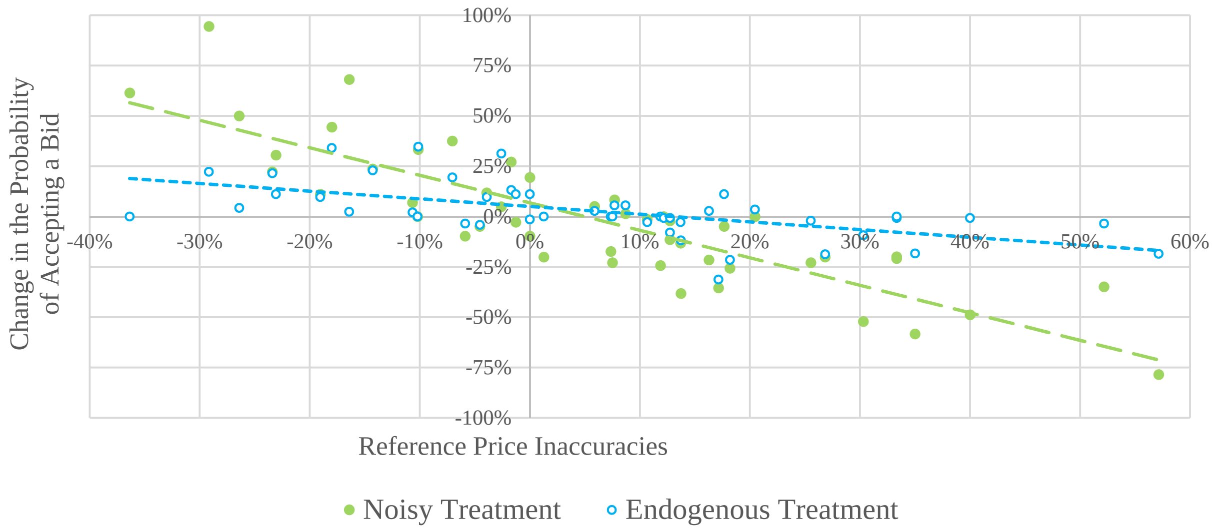

Notes: The change in the probability of accepting a bid relative to the Known treatment is plotted on the vertical axis, as a function of the reference price inaccuracy on the horizontal axis, which shows the percentage by which the reference price deviates from the actual seller value. Observations for the Endogenous treatment are indicated by hollow dots. The dashed lines represent the linear trend for each series.

Consistent with intuition, we can see in Figure 3 that there is a negative relationship between the error the seller makes when setting the reference price of a commodity and the probability that the seller accepts a bid on that commodity. That is, bids on commodities that are overvalued are less likely to be accepted, and bids on commodities that are undervalued are more likely to be accepted (as in the complete information example in Section 3). This negative relationship, however, is substantially less pronounced in the Endogenous treatment. In particular, the trend line for the Endogenous treatment (short dashed line) is flatter than the trend line for the Noisy treatment (long dashed line), a difference that is significant ($p < 0.01$). Further, the trend line the Endogenous treatment has a slope that is close to but significantly different from zero ($p < 0.05$). Thus, the seller is able to eliminate almost entirely the tendency to incorrectly accept (respectively reject) bids on commodities that are initially undervalued (respectively overvalued). This leads to our fourth result.

> ***Result 4****: Endogenous reference prices provide the seller some protection against incorrectly accepting (rejecting) securities that were initially undervalued (overvalued).*

5.4 BID DISCOUNTS

Recall that subjects in the main experiment receive the same instructions in all three treatments. In particular, they do not know how the reference prices are determined and whether or not they are accurate. Thus, bidding behavior in the first auction of each treatment should not be different. With time, however, as subjects receive different post-auction feedback, it is possible that they learn to behave differently in each treatment. If such changes in bidding behavior actually occur, then this could help explain the differences in the seller's profit across treatments that we identified in Section 5.1.

To explore whether subjects learn to behave differently across treatments, we consider the "bid discount," defined as the difference between the bidder's value and the bid submitted, expressed as a percentage of the bidder's value. For instance, a bid discount of 10% indicates that the bidder submits a bid 10% below value. Therefore, the bid discount may be interpreted as a measure of how aggressive the bidding is (i.e. a low bid discount corresponds to more aggressive bidding). We report descriptive statistics for the bid discount in column 3 of Table 2 and in Figure 4 (by auction).

Figure 4: Bid Discounts

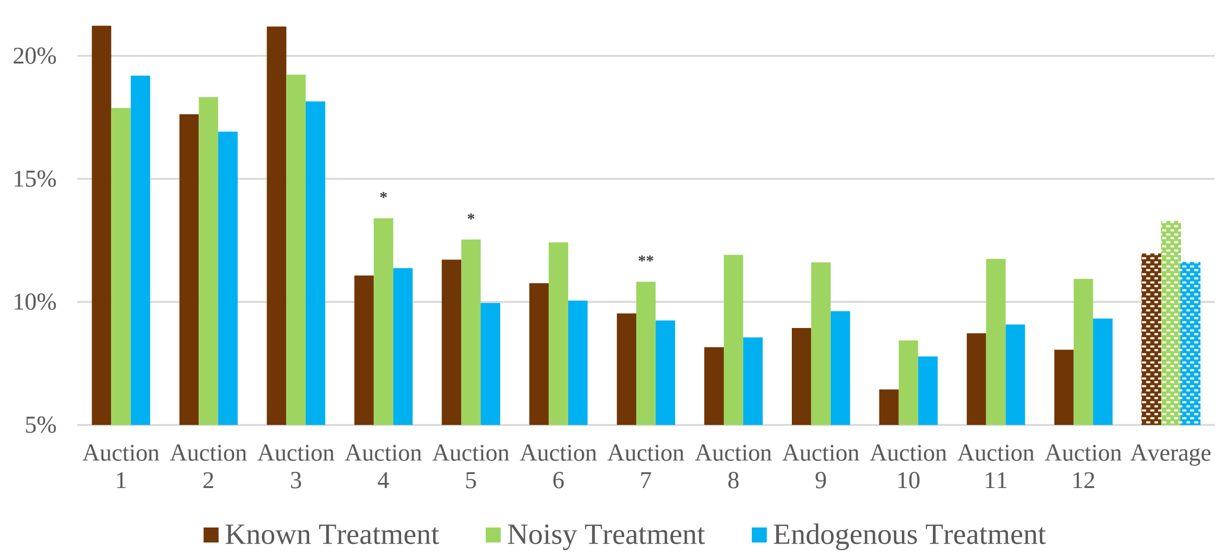

Notes: Each bar in the chart represents the average of the 6 sessions conducted for the corresponding treatment. The same random draws are used in each session and treatment, so the bars in each auction are directly comparable across treatments. The stars above a bar represents the outcome of a 2-tail permutation test of equal means between the Endogenous treatment (last bar) and the treatment corresponding to the bar under consideration. The asterisk superscripts, ***, **, and *, indicate that the null hypothesis of equal means is rejected at the 1%, 5%, and 10% significance levels.

Figure 4 confirms that the bid discounts are not statistically different across treatments in the first three auctions. This result reflects the fact that subjects in the main experiment are not made aware of any differences between treatments. After the third auction, however, it is apparent from Figure 4 that the average bid discount is systematically highest in the Noisy treatment, although the difference is most often insignificant. Thus, it appears that subjects learn to bid slightly less aggressively in the Noisy treatment. In contrast, there is no statistical evidence of a difference in bidding behavior between the Known and the Endogenous treatment in any of the 12 auctions. This result is conformed in Table 2 (column 3) where the average bid discount is similar in the Known and Endogenous treatments, and slightly (but insignificantly) higher in the Noisy treatment. This leads to our fifth result.

> ***Result 5****: Bidding behavior was slightly less aggressive in the Noisy treatment, but statistically indistinguishable in the Known and Endogenous treatments.*

The combination of results 3, 4 and 5 therefore suggests that most of the gains in seller's profit between the Noisy and the Endogenous treatments result from the improvement in the accuracy of the endogenous reference prices used to normalize the bids, not from a modification of bidding behavior.

To confirm this result, we conduct a series of counterfactual exercises in which we calculate for each auction the profit the seller would have obtained if the bids submitted in a given treatment had been scored against the reference prices used in a different treatment for the same auction. The results of this counterfactual exercise are reported in Table 3. For instance, we can see that while the actual seller's profit was on average $60.71 per auction in the Known treatment, it would have been only -$36.14 if the bids submitted in the Known treatment had been scored against the reference prices used in the Noisy treatment (see second row, first column in Table 3).[18]

Overall, Table 3 indicates that the average seller's profit does not change substantially across columns, that is, when the reference prices from a given treatment are used to score the bids submitted in different treatments. In contrast, the seller's profit is always highest when the reference prices from the Known treatment are used

[18] As noted in the caption, the seller profits in Table 3 are calculated on a per-auction basis, so numbers on the diagonal are scaled down from the totals in Table 2.

to score the bids (first row), and the seller's profit is always lowest when reference prices from the Noisy treatment are used (last row). These counterfactual exercises therefore confirm that most of the differences in seller's profit across treatments we observed in the experiment are not due to differences in bidding behavior, but instead to differences in the accuracy of the reference prices used to normalize the bids.[19]

TABLE 3—COUNTERFACTUAL COMPARISON OF SELLER'S PROFIT

		Bids Submitted in		
		Known Treatment	Noisy Treatment	Endogenous Treatment
Reference Prices used in	Known Treatment	**$60.71**	$50.08	$55.66
	Noisy Treatment	-$36.14	**-$47.48**	-41.86
	Endogenous Treatment	$23.72	$7.45	**$16.22**

Notes: In the Known treatment the seller knows its own values and reference prices are accurate. In the Noisy treatment the seller does not know its own values and must rely on noisy estimates to set the reference prices. In the Endogenous treatment, the seller has the same noisy estimates as in the Noisy treatment and sets the endogenous reference price as the average of the noisy estimate and the median of the bids submitted. Each cell represents the average seller's profit per auction calculated in a counterfactual exercise in which each auction is cleared using the bids collected in the treatment from the corresponding column are scored using the reference prices from the treatment in the corresponding row. For instance, the average seller's profit per auction obtained if the bids collected in the Known treatment (first column) had been scored against the reference prices in the Noisy treatment (second row) is -$36.14. The bold numbers on the diagonal indicate the average seller's profit actually obtained in the experiment.

6 ALTERNATIVE ENDOGENIZATION PROCESSES

Our approach to endogenize the reference prices is clearly not the only way the seller can take advantage of the information contained in the bids submitted to revise its noisy value estimates. This raises the question of whether the seller's profit can be increased further by using a different endogenization method. To address this question, we use the data collected in the Endogenous reference price treatment to conduct a series of counterfactual exercises in which we vary the manner in which reference prices are adjusted from submitted bids.

6.1 WEIGHTED AVERAGE

[19] In addition to comparing treatments with respect to the profit they generate, the seller may also be interested in possible differences in the sales revenues between the two budget constraints *Min* and *Max*. An analysis of sales revenues by session (see Appendix A) confirms that auction revenues are significantly lower for the Noisy treatments than for the Endogenous treatment, regardless of whether this procedure was explained or not ($p<0.01$ for all comparisons between Endogenous and Noisy treatments, using a 2-tailed permutation test). The endogenous correction for submitted bids seems to prevent the seller from rejecting bids that seem too high, perhaps because bidding is a little more aggressive, which could enhance sales revenue.

We begin by considering a more general approach in which the endogenous reference price for commodity j, $R_j(\bar{b}_j)$, is a *weighted* average of the seller's noisy value estimate $\hat{V}_j$ and the median bid $\bar{b}_j$: $R_j(\bar{b}_j) = w * \bar{b}_j + (1-w) * \hat{V}_j$ with $0 \leq w \leq 1$. The seller ignores the bids submitted when $w = 0$, and all weight is given to the median bid (and thus none to the seller's noisy value estimate) when $w = 1$. In the experimental sessions conducted for the Noisy and the Endogenous treatments, the weight w was set to 0 and ½, respectively. The counterfactual exercise then consists in taking the bids submitted in the Endogenous treatment and in running the allocation process again for different values of the weight w.

The validity of this counterfactual exercise rests on the assumption that the bidders' behavior would have been the same had we used a value for w different than ½ in the Endogenous treatment. This assumption finds some support in the facts that i) the subjects were unaware of how the seller determined the reference prices, and ii) the analysis in Section 5.4 revealed little difference between the subjects' bid discounts for two different values of w, i.e. $w = 0$ (the Noisy treatment) and $w = ½$ (the Endogenous treatment).

The curved solid line in Figure 5 plots the recalculated seller's profit (vertical axis) for different values of the weight w between 0 and 1 (horizontal axis). As upper and lower benchmarks, we also plot the average profit the seller obtained in the Known treatment (upper horizontal line) and in the Noisy treatment (lower horizontal line). Ignore for the moment the dotted and dashed lines which will be discussed in the next section.

Observe first on the left side of Figure 5 that the counterfactual exercise conducted with a weight $w = 0$ produces a seller profit that is similar to the seller's profit that was actually observed in the Noisy treatment (lower horizontal line). This result simply reflects the fact that, as shown in Section 5.4, bidding behavior was essentially the same in the Endogenous and in the Noisy treatments. Further, the counterfactual exercise suggests that the "optimal" weight for our data is around 0.7. Specifically, $w = 2/3$ would yield the seller a profit of \$21.7 on average per auction. This is slightly higher than the \$16.7 seller profit that was observed in the Endogenous treatment (indicated in Figure 5 by a full dot at $w = 0.5$), but this remains substantially lower than the \$60.7 we obtained in the Known treatment (upper horizontal line). Thus, using a value for w around 0.7 further mitigates the impact of noisy value estimates, but it does not solve the problem entirely.

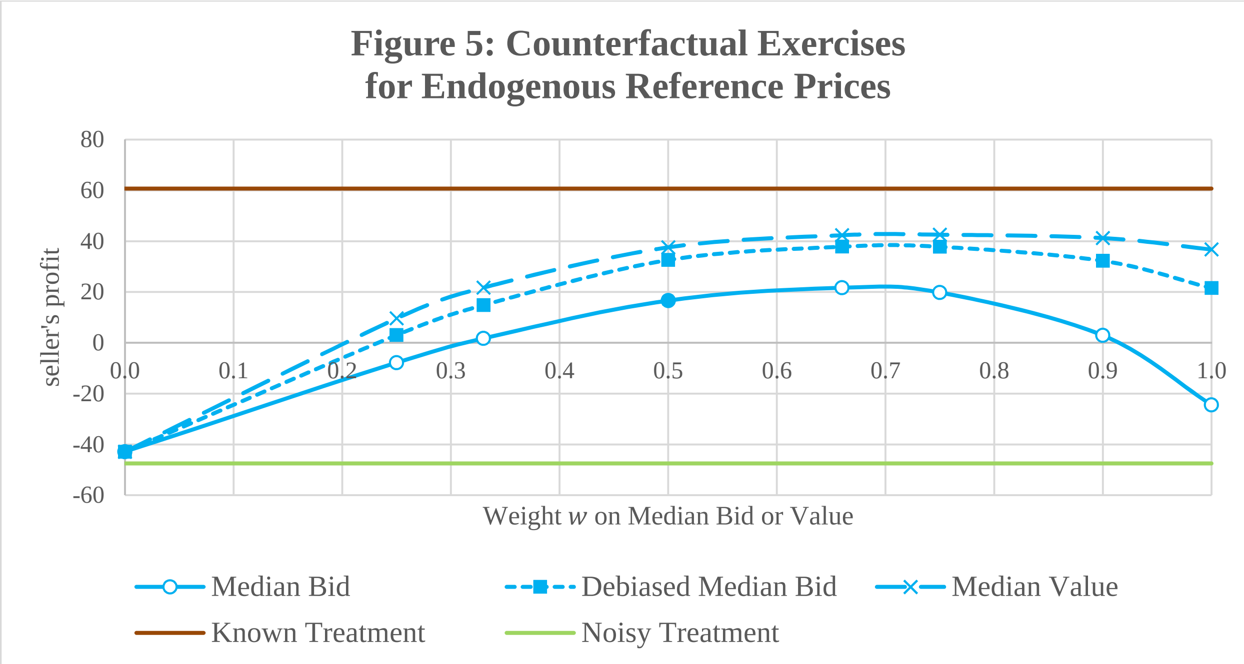

Notes: The connected hollow dots show counterfactual seller's profits when the references price is a weighted average of the seller's initial value estimate and the median bid submitted. The weight placed on the median bid submitted ranges from 0 on the left side of the horizontal axis to 1 on the right side. The full dot at $w = 1/2$ represents the seller profit actually observed in the Endogenous treatment. The lower (respectively upper) flat line indicates the seller profit actually observed in the Noisy (respectively Known) treatment corresponding to a weight of zero placed on the median bid (respectively with perfectly accurate seller value estimates). The dotted line connecting the square markers plots the recalculated seller's profit when a correction is used to debias the median bid before using the weighted average formula. The dashed line connecting the cross markers shows recalculated seller profit when reference prices are based on the median *bidder value* (instead of the median bid).

6.2 Weighted Average with Bias Correction

Although simple and easily implementable, our approach to endogenize the reference prices is in no way optimal. In particular, under the discriminatory pricing format we used in the experiment (where competition was limited to four bidders), bidders must discount their bids slightly to expect making positive profits (Table 2 shows that the bid discount exceeds 10% on average). Thus, the median bid may be a reasonable and robust estimate of the mean market value, but it is nevertheless biased downward.[20]

Can the seller somehow eliminate or at least reduce this bias? In principle, a correction would be possible if the extent to which bidders discount their bids below value was known to the seller. The bid discount, however, cannot be measured by the seller as the

[20] Our median bid approach could be considered optimal under the uniform price auction. Indeed, Klemperer (2010) (footnote 21) argues that bidding at value is a reasonable assumption at a uniform price auction when it is unlikely that any bid could have much effect on the prices paid by a bidder. Thus, with sufficient competition, bidders could be expected to bid their values, in which case the median bid submitted for a commodity is an unbiased estimate of the seller's value.

bidders' values are not observable, even after the auction. What the seller can observe after each auction, however, is its own value, which can then be compared to the median bid. In some instances, a seller who conducts multiple auctions over time may acquire sufficient knowledge through these comparisons to correct the bias in the median bid. For example, if the seller's actual values tend to be 15% above the median bids in past auctions, then the seller could multiply the median bids in the current auction by a "debiasing factor" of 1.15.

To explore the effectiveness of this debiasing method we conduct a counterfactual exercise for each experimental session of the Endogenous treatment. For auction *t=2,..,12*, we first calculate the debiasing factor δ_t equal to the average ratio of the seller's value to the median bid across all commodities sold at auctions τ *=1,..,t-1*: $\delta_t = \frac{1}{4*(t-1)} \sum_{\tau=1}^{t-1} \sum_{j=1}^{4} \left(V_{\tau j} / \bar{b}_{\tau j}\right)$. Then, the median bid for each commodity for sale at auction t is multiplied by the debiasing factor. Finally, the reference price for commodity j is set equal to the weighted average of the debiased median bid and the seller's noisy value estimate: $R_{tj}(\bar{b}_{tj}, \delta_t) = w * \left(\delta_t * \bar{b}_{tj}\right) + (1-w) * \hat{V}_{tj}$. For the first auction, we set $\delta_1 = 1$ because the seller has no information to debias the median bid.

The dotted line (with square markers) in Figure 5 plots the seller's recalculated profits when the weight given to the "debiased" median bid varies from 0 to 1. The figure indicates that endogenizing the reference prices using the debiased median bids increases the seller's profit for every value of the weight w. Using the debiased median bids and a weight between 2/3 and 3/4 yields the highest profit for the seller. In particular, a weight of $w = 2/3$ would yield the seller a profit of \$37.9 on average per auction. This is more than twice as high as the \$16.7 we obtained in the experiment for the Endogenous treatment, and 75% higher than the \$21.7 we would obtained in the previous counterfactual exercise with $w = 2/3$ and a non-debiased median bid. This counterfactual exercise therefore suggests that the seller can increase its profit substantially by using data from past auctions to debias the median bid.

6.3 PROFIT BOUND ON "OPTIMAL" ENDOGENIZATION METHOD

How does the seller's profit generated by our method compare to the profit an optimal endogenization process would generate? To answer this question, we must first characterize how the information contained in the bids can be exploited in an optimal manner by the seller to update its noisy value estimates. In principle, the best information the seller can expect to recover from the bids submitted for a commodity is an estimate of each bidder's actual value for that commodity. To do so, the seller could develop adjustments based on structural econometric estimates of the bidders' value distribution. This approach, however, requires the Nash equilibrium bidding function (or at least the first order conditions of the problem) which in a case of a multi-unit, multi-object scoring auction with endogenous reference prices is far from trivial to derive.[21]

Nevertheless, we can conduct a counterfactual exercise to gauge how much of a profit improvement such an econometric approach would provide. To do so, we assume that the seller can actually observe each bidder's values before the auction. Then, the seller can revise its reference prices for each commodity by taking a weighted average of its noisy value estimate and the median of the bidder's values. This best-case scenario provides an upper bound on the profit the seller could expect when it exploits the information contained in the bids to correct its noisy value estimates.

The dashed line (with cross markers) in Figure 5 plots the seller's recalculated profits when the weight given to the median of the bidders' values varies from 0 to 1. As expected, the seller's profit is systematically higher than the other two counterfactual exercises conducted previously in this section. Further, observe that it is optimal for the seller to give most of the weight to the median value, but the seller should not ignore its own initial noisy value estimate (which would correspond to the case $w = 1$). More importantly, note that the average seller's profit generated by the bias correction method in Section 6.2 is very close to the profit upper bound that would be obtained if the seller knew the bidder's values. In particular, the average seller profit of \$37.9 obtained for $w = 2/3$ with the bias correction method is only slightly smaller than the \$42.40 upper bound. Thus, when the seller does not know its own values, alternative methods to endogenize the reference prices are unlikely to provide major additional benefits to the seller compared to the bias correction method.

[21] An additional drawback of this econometric approach compared to our method is that it may require time to update every reference price. Having a method that can be implemented almost instantly is often a key requirement in practice. Indeed, in many auctions mere seconds can separate the closing of the auction and the announcement of the results. In financial markets, for instance, the price of a commodity allocated at the auction can vary in real time and the bidders cannot be exposed to market risk in the interim period. The need for near immediate execution also explains why "clock" or English type auctions are rare in financial markets.

7 ROBUSTNESS CHECK

The results from the main experiment suggest that using the median bid to endogenize reference prices enhances the profit of the seller when it does not know its own values. Subjects, however, did not know how the reference prices were determined. In particular, subjects in the Endogenous treatment did not know that their bids were used to adjust the reference prices. It is therefore natural to consider whether the endogenous approach would still produce higher seller's profit if bidders knew the exact process through which the noisy value estimates are transformed into final reference prices. To address this question we first conduct an additional experimental treatment.

7.1 ADDITIONAL EXPERIMENTAL TREATMENT

The additional "Explained Endogenous" treatment is identical to the Endogenous treatment with one exception: In the instructions given to the subjects at the beginning of the experiment, we explained precisely how the seller uses the median of the bids submitted for each commodity to calculate the endogenous reference price. As with other treatments, the additional treatment consists of 6 sessions, each with four different bidders and 12 auctions. The random number sequences are the same as those used before, so that treatments could be directly compared. The seller's profit and the bid discounts for this additional treatment are plotted in Figures 6 and 7.

Figure 6 indicates that, on average, the seller's profit is systematically higher (p = 0.06) in the Explained Endogenous treatment than in the Endogenous treatment. Figure 7 suggests that this result may be explained by the fact that the bid discounts are generally lower (albeit not significantly) in the Explained Endogenous treatment than in the Endogenous treatment. Thus, telling the bidders how their bids are used to calculate the reference prices does not affect bidding behavior in a manner that reduces the seller's profit. If anything, bidders are slightly more aggressive in our experiments and the seller's profit is slightly higher when the endogenization process is known to the subjects. One possibility is that competition is increased by the enhanced transparency in auction clearing procedures.

Figure 6: Seller's Profits

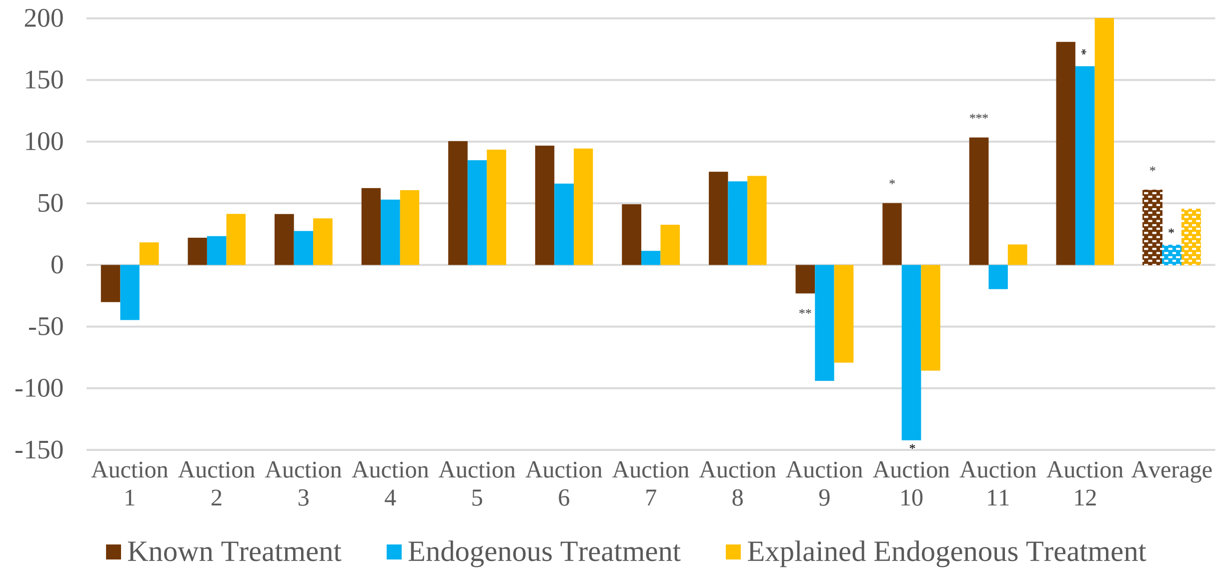

Notes: Each bar in the chart represents the average of the 6 sessions conducted for the corresponding treatment. The same random draws are used in each session and treatment, so the bars in each auction are directly comparable. The stars above a bar represent the outcomes of a 2-tail permutation test of equal means between the Explained Endogenous treatment (last bar) and the treatment corresponding to the bar under consideration. The asterisk superscripts, ***, **, *, indicate that the null hypothesis of equal means is rejected at the 1%, 5%, and 10% significance levels.

Figure 7: Bid Discounts

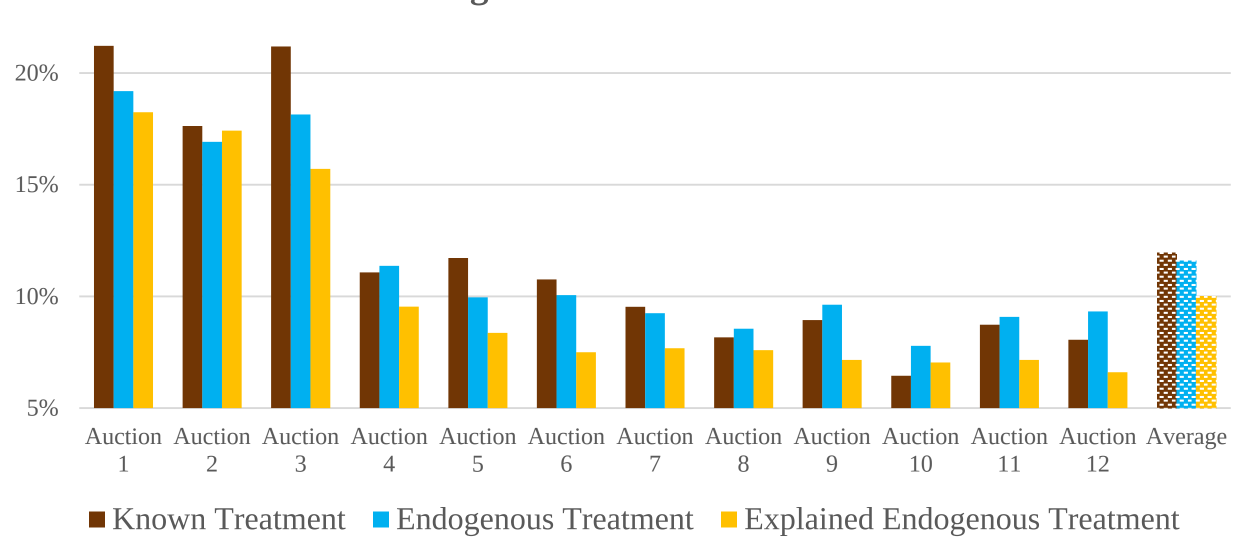

Notes: Each bar in the chart represents the average of the 6 sessions conducted for the corresponding treatment. The same random draws are used in each session and treatment, so the bars in each auction are directly comparable. The stars above a bar represent the outcomes of a 2-tail permutation test of equal means between the Explained Endogenous treatment (last bar) and the treatment corresponding to the bar under consideration. The asterisk superscripts, ***, **, *, indicate that the null hypothesis of equal means is rejected at the 1%, 5%, and 10% significance levels.

7.1 POSSIBLE BIDDING MANIPULATION

The experiment we conducted, including the robustness check treatment, provided no evidence that subjects adjusted their bidding strategies in a way detrimental to the seller. In principle, however, bidding manipulation cannot be excluded when the endogenization process is known. In particular, a bidder may have an incentive to submit a "phantom bid," i.e. a bid for a large quantity at a very low price (e.g. at the reserve price), with the sole objective of lowering a commodity's median bid. By doing so, the bidder may lower the endogenous reference price and make another one of his bid look more attractive to the seller. Note also that submitting a phantom bid is essentially cost-free to the bidder because such low bids have little chance of being accepted.[22]

So, how can the seller protect itself against such phantom bids? One way to do so is simply not to disclose how reference prices are set. If bidders are not aware that their bids are used to determine reference prices, then they do not know that phantom bids can be advantageous to them.[23] Additional measures to mitigate the incidence of phantom bids include setting informed and reasonable reserve prices, limiting the total number of bids (i.e. price-quantity pairs) a bidder can submit per commodity, and limiting the amount of a commodity a bidder can bid for. To remove any incentive to submit phantom bids, the seller could also modify the endogenization process slightly. For instance, instead of using the median of *all* the bids submitted for a commodity, the seller could use the median of the *highest* bid submitted by each bidder for that commodity. This way, low phantom bids would not be taken into consideration. Because bidding above value is always dominated in a private value auction, the median of the highest bids submitted should still produce a good, albeit biased, estimate of the seller's value.[24]

To gauge how this modified endogenization process would have performed in our experiment, we conducted a final counterfactual exercise using the bids collected in the

[22] In contrast, it appears unlikely that a bidder would try to influence the median bid by submitting a higher than normal bid (e.g. above the bidder's value) because such high bids would have a high probability of being accepted, which would be costly to the bidder.

[23] As mentioned earlier, the allocation process in many auctions is often not entirely explicit. Thus, not disclosing the endogenization process would not be uncommon in practice. Further, the seller's incentive for secrecy would explain why this endogenization process would not be publicly known even if it had been used.

[24] The bias in the median of the highest bids might be less than using the median of all bids, but any remaining bias could be reduced with the correction method described in section 6.1. Note that this modified endogenization method would also provide a good value estimate in a pure common value auction.

explained endogenous treatment. Interestingly, we find that calculating endogenous reference prices using only the highest bid submitted by each bidder for a commodity produces a seller's profit that is slightly (but not significantly) higher on average ($48.56 in the counterfactual exercise versus $47.07 in the endogenous explained treatment), but substantially more volatile from auction to auction (the standard deviation is $121.24 in the counterfactual exercise versus $89.49 in the endogenous explained treatment). These results most likely reflect the fact that using only the highest bid submitted by each bidder for a commodity produces an estimate of the seller's value that is slightly less biased (because it relies on the bids closest to the bidders' values), but more variable (because it is calculated with fewer data points). Thus, the counterfactual exercise illustrates how a slight modification of the endogenization process can remove incentives for bidding manipulations without reducing average seller's profit.

We conclude with section with a short discussion of possible collusion. Collusion is a legitimate concern for any auction, and more specifically for multi-unit auctions (Goswami et al. 1996, Kremer and Nyborg 2004, Sade et al. 2006). As explained in Armantier et al. (2013), one of the key benefits of a reference price auction is to promote competition by creating thick markets in which bidders for different commodities compete against each other.[25] Because of this cross-commodity competition, a well-designed reference price auction should attract a large number of different bidders, thereby making coordination more difficult. Thus, although not immune to collusion, an endogenous reference price auction is less susceptible to the problem.

8 CONCLUSION

When an auctioneer decides which bids to accept at a multi-object, multi-unit auction, it must be able to compare bids across different commodities. If the auctioneer knows its own value for each commodity (as in the standard private values model), then the "reference price" allocation mechanism (Armantier et al. 2013) enables the auctioneer to accept the bids with

[25] To illustrate how a reference price auction promotes competition, consider a bidder who is willing to buy 2 units of A and another bidder who is willing to buy 2 units of B. The seller is willing to sell up to 2 units of A or B. The seller earns low profit by conducting a separate auction for a unit of each object because each auction has a single bidder. In contrast, a reference price auction promotes competition by having the two bidders compete against each other.

the best relative values. This paper focused on the case where the auctioneer must rely on inaccurate value estimates. We proposed a generalization of the reference price auction in which the auctioneer reduces value inaccuracies by setting references prices endogenously using information contained in the bids submitted. Specifically, the reference price for a commodity was set midway between the auctioneer's inaccurate value estimate and the median of the bids submitted for that commodity. We conducted an experiment which confirms that such an endogenous approach increases the auctioneer's profit significantly. A counterfactual analysis suggests that using data from past auctions to set reference prices can help the auctioneer raise profit near the upper bound from any endogenous method.

These results have important practical implications because many important auctions, especially in finance, are multi-object, multi-unit and involve a budget constrained auctioneer who may have imperfect value estimates.[26] Among the various examples presented in the introduction, consider for instance Treasury buyback auctions. In 2012, the General Accounting Office (GAO) recommended to the U.S. Treasury the use of regular buyback auctions as an effective debt management tool, both in time of surplus and deficit.[27] In October 2014, the Treasury started conducting small-value buyback auctions to ensure operational readiness. Although the GAO recommends auctions that include multiple bonds and emphasizes the importance of accepting the most competitive bids, it does not discuss how the Treasury should "cherry pick" among bids for different bonds to achieve this objective.[28] A reference price auction would provide a simple yet effective solution to this problem. Further, because Treasury bonds have a strong common value component, and because the Treasury is not an active trader on the secondary market, the value of each bond is not likely to be known perfectly by the Treasury at the time of the auction. Thus, endogenizing the reference prices using an approach similar to the one proposed in this paper should benefit the Treasury as it would help accept the bids with the highest relative values.

[26] Hendershott and Madhavan (2015) report that financial markets are increasingly transitioning from "voice to electronics," i.e. from over-the-counter trading to electronic auctions.
[27] See also similar recommendations by the IMF for emerging countries (Medeiros, Polan and Ramlogan 2007), as well as the arguments in favor of regular Treasury buyback auctions in Garbade and Rutherford (2007).
28 The importance of a systematic "cherry picking" allocation mechanism is illustrated by Garbade and Rutherford (2007) who described the settlement of the largest auction in the 2001 U.S. Treasury buyback as "unwieldy" because it attracted a large number of bids for 26 different bonds that took time to process and compare during the five minutes that separated the end of the auction and the announcement of the results.

We believe that the contribution of this paper is not limited to the manner in which we propose to endogenize reference prices. More generally, it is in the recognition that, in some instances, the auctioneer can use the information contained in the bids submitted to its own benefit. Such strategic behavior on the part of the auctioneer is not uncommon in practice. In particular, several Treasuries (e.g. in Switzerland or Finland) adjust the quantity of bonds issued depending on the demand expressed by bidders at the auction.[29] This practice has been shown to benefit the auctioneer (Back and Zender 2001, McAdams 2007).[30] While the literature has focused on the bidders' incentives to acquire information (Milgrom and Weber 1982), we considered a situation in which the auctioneer also has an incentive to acquire information and we proposed a simple, cost-free mechanism to do so effectively.

Our results are also informative for single-object auctions. For instance, in the canonical model of Myerson (1981) the seller can raise profit by pre-specifying an optimal reserve price that exceeds its own value. In practice, reserve prices are commonly used but they are often kept secret, which has been somewhat of a "puzzle" for auction theorists (Bajari and Hortaçsu 2003, Jehiel and Lamy 2015). In the environment considered in this paper, the auctioneer's value is not known but it is correlated with the bidders' values, and the auctioneer may be able to improve the accuracy of its value estimate by using the bids submitted. In that case, the seller may have an incentive to keep the reserve price secret and to adjust it endogenously using the bids submitted. In fact, such practice may be relatively common. Indeed, some real world auctions clearly allow the seller to adjust its reference price during the auction,[31] while most auctions with secret reserve prices do not explicitly prevent the seller from doing so.

References

Abbink K., Brandts J. and P. Pezanis-Christou (2006) "Auctions for government securities: A laboratory comparison," *Journal of Economic Behavior and Organization*, 61(2), 284-303.

Armantier O. (1999) "Three essays on the empirical application of game theoretic models," Ph.D. Dissertation, University of Pittsburgh.

[29] As in our main experiment, the bidders in these auctions are not informed about how their demands affect the quantity of bonds issued.

[30] Brenner et al. (2009) report that 30 of the 48 Treasuries they surveyed have the ability to adjust the quantity of debt issued after opening the bids. See also Umlauf (1993).

[31] For instance, eBay allows a seller to lower or remove its secret reserve price during the auction.

Armantier O., Holt C. and C. Plott (2013) “A Procurement auction for toxic assets with asymmetric information,” *American Economic Journal: Microeconomics*, 5(4), 142-162.

Armantier O. and J. Sporn (2013) “Auctions implemented by the Federal Reserve Bank of New York during the Great Recession,” *Federal Reserve Bank of NY Staff Reports*, 635.

Asker J. and B. Cantillon (2008) “Properties of scoring auctions,” *RAND Journal of Economics*, 39(1), 69-85.

Back K. and J. Zender (1993) “Auctions of divisible goods: On the rationale for the Treasury experiment,” *Review of Financial Studies*, 6, 733-764.

Back K. and J. Zender (2001) “Auctions of divisible goods with endogenous supply,” *Economics Letters*, 73, 29-34.

Bajari P. and A. Hortaçsu (2003) “The winner’s curse, reserve prices, and endogenous entry: Empirical insights from eBay auctions,” *RAND Journal of Economics*, 34(2), 329–55.

Baird D. (1993) “Revisiting auctions in Chapter 11,” *The Journal of Law and Economics*, 36(1), 633-53.

Brenner M., Galai D. and O. Sade (2009) “Sovereign debt auctions: Uniform or discriminatory?” *Journal of Monetary Economics*, 56, 267-274.

Cason T., Gangadharan L. and C. Duke (2003) “A laboratory study of auctions for reducing non-point source pollution,” *Journal of Environmental Economics*, 46, 446-471.

Chen-Ritzo C., Harrison T., Kwasnica A. and D. Thomas (2005) “Better, faster, cheaper: An experimental analysis of a multiattribute reverse auction mechanism with restricted information feedback,” *Management Science*, 56, 1753-1762.

Cox J., Smith V. and J. Walker (1984) “Theory and behavior of multiple unit discriminative auctions,” *Journal of Finance*, 39, 983-1010.

Frost T., Govier N. and T. Horn (2015) “Innovations in the Bank's provision of liquidity insurance via indexed long-term repo operations,” *Bank of England Quarterly Bulletin*.

Garbade K. and M. Rutherford (2007) “Buybacks in Treasury cash and debt management,” *Federal Reserve Bank of New York Staff Reports*, 304.

Goeree J. and C. Holt (2008) “Standardized uniform-price reverse auctions: A multi-market auction approach for the Treasury’s rescue plan,” draft proposal submitted to Treasury, October 9, 2008.

Goswami G., Noe T. and M. Rebello. (1996) “Collusion in uniform-price auctions: experimental evidence and implications for Treasury auctions,” *The Review of Financial Studies*, 9(3), 757–785.

Han B., Longstaff F. and C. Merril (2007) “The U.S. treasury buyback auctions: The cost of retiring illiquid bonds,” *The Journal of Finance*, 62(6): 2673-93.

Hendershott T. and A. Madhavan (2015), "Click or Call? Auction versus Search in the Over-the-Counter Market," *The Journal of Finance*, 70(1), 419-447.

Hortaçsu A. and D. McAdams (2010) "Mechanism choice and strategic bidding in divisible good auctions: An empirical analysis of the Turkish Treasury auction market," *Journal of Political Economy*, 118(5), 833–865.

Jehiel P. and L. Lamy (2015) "On absolute auctions and secret reserve prices," *RAND Journal of Economics*, 46(2), 241–270.

Klemperer P. (2010) "The product-mix auction: A new auction for differentiated goods," *Journal of the European Economic Association*, 8(2-3), 526–536.

Kremer I. and K. Nyborg (2004) "Underpricing and market power in uniform price auctions," *Review of Financial Studies*, 35(1), 147-159.

Kwasnica A. and K. Sherstyuk (2013) "Multiunit Auctions," *Journal of Economic Surveys*, 27(3), 461-490.

McAdams D. (2007) "Adjustable supply in uniform price auctions: Non-commitment as a strategic tool," *Economics Letters*, 95(1), 48–53.

Medeiros C., Polan M and P. Ramlogan (2007) "A primer on sovereign debt buybacks and swaps," IMF Working Paper No. 07/58.

Milgrom P. and R. Weber (1982) "A theory of auctions and competitive bidding" *Econometrica*, 50(5), 1089–1122.

Morales-Camargo E., Sade O., Schnitzlein C. and J. Zender (2013) "Divisible good auctions with asymmetric information: an experimental investigation," *Journal of Financial and Quantitative Analysis*, 48(4), 1271-1300.

Myerson R. (1981) "Optimal Auction Design," *Mathematics Operations Research*, 6, 58-73.

Nyborg K. and S. Sundaresan (1996) "Discriminatory versus uniform Treasury auctions: evidence from when-issued transactions," *Journal of Financial Economics*, 42, 63-105.

Sade O., Schnitzlein C. and J. Zender (2006) "Competition and cooperation in divisible good auctions: An experimental examination," *Review of Financial Studies*, 19(1), 195–235.

Stoneham G., Chundhri V., Ha A. and L. Strappazzon (2002) "Victoria's bush tender trial: A cost sharing approach to biodiversity," discussion paper, presented at the Inaugural Sheep and Wool Industry Conference, Hamilton, Victoria, Australia.

Umlauf S. (1993) "An empirical study of the Mexican T-bill auction," *Journal of Financial Economics*, 33, 313–340.

Appendix A. Summary Performance Measures by Auction Sequence

Session Name	Reference Price Treatment	Sales Revenue	Seller Profit	Reference Price Inaccuracy	Bid Discount Percentage
(tbae7)	Known 1	18751	1087.73	__	9.27
(tbae8)	Known 2	17953	905.81	__	10.85
(tbae9)	Known 3	18638	861.25	__	10.59
(tbae16)	Known 4	16533	503.25	__	13.17
(tbae17)	Known 5	14839	77.37	__	17.17
(tbae18)	Known 6	18146	935.89	__	10.66
	Average	**17477**	**728.55**	**__**	**11.95**
(tbae4)	Noisy 1	15776	-1189.27	17.64%	18.58
(tbae5)	Noisy 2	17530	-467.32	17.64%	12.92
(tbae6)	Noisy 3	17498	-299.96	17.64%	12.69
(tbae13)	Noisy 4	16861	-235.12	17.64%	10.13
(tbae14)	Noisy 5	16673	-694.48	17.64%	13.17
(tbae15)	Noisy 6	16580	-532.71	17.64%	12.13
	Average	**16820**	**-569.81**	17.64%	**13.27**
(tbae1)	Endogenous 1	19626	521.82	5.78%	9.28
(tbae2)	Endogenous 2	19611	418.94	6.05%	10.48
(tbae3)	Endogenous 3	19387	344.45	5.98%	11.06
(tbae10)	Endogenous 4	18411	-251.47	6.81%	15.25
(tbae11)	Endogenous 5	19820	273.5	5.59%	10.37
(tbae12)	Endogenous 6	19410	-139.75	4.40%	13.19
	Average	**19378**	**194.58**	**5.77%**	**11.60**
(tbae19)	Explained Endogenous 1	20386	759.76	6.23%	7.39
(tbae20)	Explained Endogenous 2	19144	175.62	5.66%	11.14
(tbae21)	Explained Endogenous 3	19689	469.14	5.45%	9.56
(tbae22)	Explained Endogenous 4	20838	711.32	6.88%	12.09
(tbae23)	Explained Endogenous 5	21748	917.32	5.11%	6.85
(tbae24)	Explained Endogenous 6	19417	356.34	4.86%	10.72
	Average	**20204**	**564.91**	**5.70%**	**9.62**