ORIGINAL PAPER

Endogenous reference price auctions for a diverse set of commodities: an experimental analysis

Olivier Armantier[1] · Charles A. Holt[2]

Abstract

This paper is concerned with multi-object, multi-unit auctions with a budget constrained auctioneer who has noisy value estimates for each object. We propose a new allocation mechanism, the *endogenous reference price auction*, with two key features. First, bids are normalized across objects using "reference prices." Second, reference prices are set endogenously using information extracted from the bids submitted. We report on an experiment showing that a simple endogenous process mitigates value inaccuracies and improves three performance measures: the seller's profit, allocative efficiency and total surplus. These results have important implications for large auctions used in practice.

Keywords Auction design · Laboratory experiment · Allocation mechanism

JEL Classification G10 · D44 · C92 · E58

1 Introduction

This paper is concerned with the allocation mechanism in an original auction environment: a multi-object, multi-unit auction with a budget constrained and imperfectly informed auctioneer. Specifically, we study how an auctioneer who receives bids to purchase (or sell) shares of different commodities should decide which bids to accept conditional on a budget constraint (e.g. a specific amount to raise or to spend) when he does not know precisely his own value for each commodity. This allocation problem is of great importance as many large auctions share these features.

✉ Olivier Armantier
olast1@gmail.com

1   Federal Reserve Bank of New York, New York, NY, USA

2   The University of Virginia, Charlottesville, Virginia, USA

⚛ Springer

Practical examples include: (i) The Treasury *buyback* auctions in which the Treasury sets a fixed notional amount of different bonds to be purchased simultaneously at auction[1]; (ii) The Quantitative Easing (QE) auctions in which a pre-specified amount of central bank reserves is used to acquire simultaneously different government bonds[2]; (iii) The liquidity auctions conducted by central banks in which a set amount of liquidity is allocated across financial institutions against different types of collateral[3]; (iv) The open market operations conducted by central banks in which a specific amount of reserves is drained (or injected) by purchasing (or selling) various types of securities simultaneously at auction[4]; (v) The 2008 TARP program which initially called for the use of auctions to purchase $700 billion worth of different "toxic assets"[5]; (vi) The "Bids Wanted in Competition" (BWIC) auctions in which a financial institution raises a specific amount of liquidity by circulating a list of different bonds for which it is willing to accept offers. In each of these examples, the auctioneer has a liquidity obligation that is met by buying or selling units of different commodities whose values the auctioneer may not know precisely ex-ante.

As documented in Sect. 2, the literature provides little guidance on how the auctioneer in these practical examples should compare and accept bids across different commodities. One exception is Armantier et al. (2013) (AHP hereafter) who proposed a relatively simple approach, a "reference price" auction. As explained in Sect. 3, a reference price auction is an allocation mechanism in which the auctioneer first normalizes the bids across commodity by dividing them by a commodity specific value estimate called a "reference price." Then, the highest normalized bids are accepted until the auctioneer's 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. Further, the approach promotes competition by letting bidders

---

[1] Akcadag et al. (2015) report that 88% of OECD countries conduct such buyback (or swap) auctions. The U.S. Treasury ran its largest buyback operation in 2000–2002 when $67.5 billion worth of various bonds were bought back in 45 auctions (Han, Longstaff and Merril, 2007). In each of these auctions, the Treasury set a fixed notional amount to be purchased and could cherry pick among the bids submitted for up to 26 different bonds (Garbade, 2015). Although smaller in size, the U.S. Treasury is still conducting small-value buyback auctions on a regular basis.

[2] In response to the Covid-19 pandemic, the Federal Reserve, the European Central Bank, and the Bank of England launched a $700 billion, €1.35 trillion and £745 billion QE program, respectively. QE auctions were also used by the Federal Reserve between 2009 and 2011 to purchase $1.19 trillion worth of various Treasury bonds (Song and Zhu, 2018) and by the Bank of England to purchase £435 billion worth of domestic government securities (Breedon, 2018).

[3] Notably, the Bank of England conducts liquidity auctions in which three separate classes of collateral (from less to more risky) are considered simultaneously (see Klemperer, 2010, or Frost, Govier and Horn, 2015). The Reserve Bank of Australia, the Bank of Japan and the Bank of Canada also conduct regular liquidity auctions. Other central banks conducted liquidity auctions during the great recession. In particular, the Federal Reserve conducted 60 auctions between 2007 and 2010 under the Term Auction Facility (TAF) program in which $3.8 trillion in loans were allocated to 429 different financial institutions against collateral of varying riskiness (Armantier and Sporn, 2013).

[4] The Federal Reserve has generally implemented monetary policy using "temporary open market operations" which consist of 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.

[5] These proposed auctions are described and evaluated in Armantier et al. (2013).

compete across commodities on a "leveled playing field" (i.e. controlling for differences in commodities' values). Of course, the performance of the reference price auction is contingent on the quality of the auctioneer's value estimates. The laboratory experiments conducted by AHP in a common value environment confirmed that reference price auctions perform well when the auctioneer knows his own value for each commodity. AHP however, also found that auction performance can deteriorate substantially when the auctioneer is imperfectly informed and must rely on noisy value estimates to set reference prices. So how can auction performance be improved in such cases?

In this paper, we attempt to answer this question by proposing and testing experimentally a new mechanism, the *endogenous* reference price auction. The mechanism rests on the idea that if the auctioneer's values are correlated with bidders' own values, then the bids submitted are also informative about the auctioneer's values. Thus, the auctioneer could use the bids to sharpen value estimates and set more accurate endogenous reference prices. To illustrate how such an approach may improve auction performance, we consider a simple process in which each endogenous reference price is set midway between the seller's initial noisy value estimate and the median of the bids submitted for the corresponding commodity. To gauge how this process would perform in practice, we conduct a policy-oriented experiment with financially motivated human subjects. Our results indicate that even a simple endogenization of the reference prices can help the auctioneer mitigate value inaccuracies and significantly improve three performance measures: the auctioneer's profit, allocative efficiency and total surplus.[6]

This paper contributes to the literature in at least three ways. First, we consider an auction environment (a multi-unit, multi-object auction with a budget constrained and imperfectly informed auctioneer) with several important practical applications which has received little attention in the literature. Second, we propose a new allocation mechanism, the endogenous reference price auction, applicable in general settings (e.g. common or private values, uniform or discriminatory auctions). Third, we provide experimental evidence that even a simple endogenization process can improve auction performance significantly.

The related literature is summarized in Sect. 2. The environment and the endogenous reference price auction are described in Sect. 3. We illustrate the approach with a complete information example in Sect. 4. Section 5 describes the experimental treatments and theoretical predictions. Auction performance is compared across treatments in Sect. 6. Section 7 explores the channels through which the endogenization process improves auction performance. Alternative endogenous procedures are considered in Sect. 8. Finally, we examine how making the endogenous process explicit can affect auction performance in Sect. 9.

---

[6] Although there is evidence suggesting that treatment effects generally carry over from typical lab subjects to professionals (see e.g. Fréchette, 2016; Snowberg and Yariv, 2018; or Dyer and Kagel, 1996 for an auction example), we are cognizant of external validity concerns. So, while we provide some evidence that an endogenous reference price auction can mitigate noisy seller's estimates, this result will need to be confirmed with using data.

2 Related literature

This paper is an extension of AHP in which we propose a new endogenous allocation method to improve the performance of AHP's reference price auction when the auctioneer is imperfectly informed about the value of the objects for sale at the auction. The experimental setup used to study the reference price auction is also different in the two papers. AHP consider a reverse (procurement) auction, a common value setting and a uniform price (single-price) mechanism in which the auctioneer is a buyer, whereas the experiment in this paper is based on an affiliated private values, discriminatory (pay-your-bid) auction in which the auctioneer is a seller.

Although the literatures on multi-unit auctions and on multi-object auctions are quite extensive,[7] few papers have studied multi-object, multi-unit auctions in which shares of different objects are sold. In addition to AHP, one exception is Klemperer (2010) who proposed the *product-mix* auction. In their basic forms, the reference price and product-mix auctions are closely related. In particular, it is easy to show that they are equivalent in the environment considered in this paper where the auctioneer's value for each object is fixed (i.e. it does not change with what is sold at the auction). The key difference with Klemperer (2010) is that he assumes that the auctioneer knows his own value for each object, whereas we consider an imperfectly informed auctioneer who can use the bids submitted to improve auction performance. The product-mix auction, however, would also underperform if the seller had noisy value estimates. Thus, exploring how to mitigate value inaccuracies endogenously is relevant for the product-mix auction as well.

The paper is also related to the literature on scoring auctions. In particular, Asker and Cantillon (2008) conduct a theoretical analysis of linear scoring auctions, whereas Cason et al. (2003), and Chen-Ritzo et al. (2005) use experimental methods to investigate the properties of single-unit scoring auctions. What sets the reference price auction apart from scoring auctions is its purpose. The main objective of standard scoring auctions is to summarize the multidimensional attributes of a bid (e.g. price, quality, time of completion) into an easily comparable unidimensional score. In contrast, as explained in Sect. 3, the purpose of a reference price auction is to control for differences in commodities' values so as to create thick markets in which bidders interested in different commodities compete on a "leveled playing field."

More generally, this paper builds on the literature that uses experimental methods to testbed mechanisms for which precise theoretical predictions are not always available. Multi-object and multi-unit auctions are classic examples of complex problems and the experiments of (e.g.) Cox et al. (1984), Cason (1995), Kagel and Levin (2001) or Sade et al. (2006) have greatly contributed to improving our understanding of these mechanisms. More generally, Roth (2016) reviews how testbed experiments have been used in market design.

---

[7] For reviews, see Kwasnica and Sherstyuk (2013), Kagel and Levin (2016), or Bichler and Goeree (2017).

Finally, the paper is related to the literature that studies how a strategic auctioneer can improve auction performance. For instance, Myerson (1981) studies how the auctioneer can set reserve prices optimally to raise profits, while Milgrom and Weber (1982) explore the auctioneer's incentives to conceal or disclose information to bidders. Some papers have studied other forms of endogenous adjustments to the allocation process based on the bids submitted. For instance, Brenner et al. (2009) report that 30 of the 48 national treasuries they surveyed can adjust the quantity of bonds they issue at Treasury auctions based on the aggregate amount bid, a practice that can benefit the auctioneer (Back & Zender, 2001; McAdams, 2007). Similarly, the amount of liquidity allocated at the Bank of England liquidity auctions is adjusted based on the aggregate amount bid (Frost et al., 2015). In this paper, we consider a new situation in which an imperfectly informed auctioneer wants to improve auction performance by extracting information from the bids submitted at auction, and we propose a simple, cost-free mechanism to do so effectively.

3 The auction environment

In this section, we consider a general multi-object, multi-unit auction environment. At this stage, we do not restrict attention to common or private values, we do not impose a pricing mechanism (uniform or discriminatory), and we leave the seller's objective function unspecified (as we shall see, the seller may consider profits, allocative efficiency and total surplus).[8] The discussion below pertains to the case in which the auctioneer is the seller, but our approach applies equally to procurements in which the auctioneer is the buyer.

Consider a seller with a portfolio of J commodities. The seller owns k_j units of commodity $j = 1, \ldots, J$. The value of a unit of j to the seller is V_j. In the canonical auction model (e.g. Milgrom & Weber, 1982), the seller knows V_j before the auction. In this paper, we focus on the more general case in which the seller does not know his values V_j and has to rely on noisy signals $\widehat{V}_j$. For instance, the signal may be the commodity's market price at the time of the auction. This signal, however, may not fully capture the fundamental value of the commodity (e.g. when markets are impaired), or it may be different from the value the auctioneer privately assigns to the commodity. The seller conducts an auction in which he is willing to sell any part of his portfolio to meet a liquidity obligation. The seller's revenue objective is characterized by lower and upper targets, *Min* and *Max*. The seller has an exact revenue objective when *Min* = *Max*, and he can maintain a balanced portfolio by adding commodity specific budget constraints, Min_j and Max_j. Each bidder can submit multiple bids for each commodity, where a bid is a price for a unit of a commodity.[9]

---

[8] The U.S. Treasury aims to optimize both profits *and* efficiency when conducting buyback auctions (Han et al., 2007).

[9] Alternatively, a bid can be defined as a price per unit and a number of units. Note also that we do not have to specify the bidders' signals at this point because the approach so far applies equally to common and private values auctions.

In this environment, how can the seller compare bids across commodities and decide which bids to accept? The simplest allocation mechanism is probably a *combined auction* in which all bids are pooled together and the highest bids are accepted first irrespective of commodity until the seller's revenue objective is met. This approach may promote competition by letting bidders compete across commodities, but adverse selection is often severe in combined auctions (see Varian, 2008 and the example in Sect. 4). A natural alternative is a set of *simultaneous auctions*, one per commodity, each with its own revenue target. This approach effectively deals with value heterogeneity across commodities, but it does not promote cross-commodity competition. As a result, performance suffers when competition for a commodity is low due to a thin market.

AHP proposed an alternative allocation mechanism, the reference price auction, under which the seller accepts the bids deemed to have the best relative values across commodities. The reference price auction consists of three steps. First, 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. A normalized bid strictly above 1 is favorable from the seller's perspective (because it exceeds his value estimate), but a normalized bid below 1 is unfavorable. Second, 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 all the bids are accepted or until the minimum revenue target *Min* is reached. Third, the seller keeps accepting bids in decreasing order of normalized bids as long as (i) the bids are favorable (i.e. with a normalized bid strictly above 1) and (ii) the maximum revenue target *Max* is not reached. The seller stops accepting bids when one of those two conditions is not satisfied.[10]

In principle, the reference price auction combines the benefits of the *simultaneous auctions* (i.e. controlling for value heterogeneity across commodities) and the *combined auction* (i.e. promoting cross-commodity competition). The experiment conducted by AHP (in a common-value uniform-price environment) confirms that the reference price auction performs well when the seller knows his own value for each commodity before the auction, in which case reference prices can be set accurately (i.e. $R_j = V_j$). However, AHP also found that auction performance (i.e. the seller's profit, efficiency) suffers when the seller does not know V_j at the time of the auction and has to rely on the noisy value estimate $\widehat{V}_j$ to set the reference price (i.e. $R_j = \widehat{V}_j$).

In this paper, we propose an approach aimed at mitigating the impact of noisy value estimates. To be clear, the objective is not to find an optimal mechanism. Instead, we test whether a simple adaptive process can reduce reference price inaccuracies and improve auction performance. Thus, consistent with Roth (2002) and

---

[10] Like the *combined* and *simultaneous* auctions, the reference price auction needs to be paired with a pricing mechanism such as the standard discriminatory or uniform-price mechanism. Which pricing mechanism dominates from the auctioneer's perspective needs to be determined on a case-by-case basis (Hortaçsu and McAdams, 2010).

Bolton and Ockenfels (2012), we place ourselves more in the position of an engineer who establishes a proof of concept, than an auction theorist who attempts to identify an optimal design. The premise for our approach is that if the seller's and bidders' values are correlated, then the bids submitted for a commodity should provide the seller with additional information about his own value. The seller can then exploit this information to revise his initial value estimate $\widehat{V}_j$ and set more accurate *endogenous* reference prices.

To gauge whether endogenous reference prices can improve auction performance we complete the specification of the auction environment: In the remainder, bidders' have affiliated private values drawn around the seller's actual value for each commodity, and we evaluate a simple approach in which each *endogenous reference price* is set midway between the seller's initial value estimate $\widehat{V}_j$ and the median of the bids submitted for that commodity $\overline{b}_j$: $R_j\left(\overline{b}_j\right) = 0.5\widehat{V}_j + 0.5\overline{b}_j$.

Before moving on, it is useful to reconsider the practical examples discussed in introduction. Which of the allocation mechanisms presented so far are actually used in these auctions? This question is difficult to answer because the information made available to bidders is often vague or inexistent. For instance, the rules regulating the U.S. Treasury buyback auctions make no mention of how bids are accepted.[11] Similarly, BWICs announcements do not specify how the seller compares bids across securities. In the open market operations and QE auctions conducted by the Federal Reserve, the allocation process is not fully disclosed but it is similar to the reference price auction.[12] For the 2008 TARP program, the Treasury initially considered combined and simultaneous auctions, but then settled on a reference price design.[13] Finally, the liquidity auctions conducted in the U.S. used a combined auction, while those in the UK follow the product-mix format. This discussion raises one question: Why do many auctioneers keep the allocation mechanism partly or totally secret? The reason most often mentioned by practitioners is to avoid bidder manipulations.[14] Examples of such possible manipulations and the benefits of concealing the allocation process are discussed in Sect. 9.

---

[11] See Department of Treasury "Marketable Treasury Securities Redemption Operations; Final Rule." Garbade (2015) also notes (p. 100): "The Treasury does not disclose the criteria that it uses in accepting one offer in lieu of another."

[12] For QE auctions, the rules state: "Offers will be evaluated based on their proximity to prevailing market prices at the close of the auction as well as on measures of relative value. Relative value measures are calculated using the Federal Reserve Bank of New York's proprietary model." https://www.newyorkfed.org/markets/longertermtreas_faq.html

[13] As explained in AHP, the Treasury ultimately decided to abandon auctions in favor of capital injections.

[14] Many other practical auctions do not fully disclose the allocation mechanism. For instance, bidders in the Treasury and liquidity auctions discussed in Sect. 2 are not told how their bids are used to set endogenously the amount of bonds or the amount of liquidity issued. The auctioneer may have other reasons to conceal information. For instance, the seller can increase his profit by keeping the actual number of bidders secret (Dyer, Kagel and Levin, 1989).

4 Illustrative example

To start, we consider a complete information example that illustrates how endogenous reference prices can improve auction performance when the seller is imperfectly informed about his values. Consider a seller who is willing to sell at auction up to three units of commodity A and three units of commodity B to meet a revenue objective between $Min = 125$ and $Max = 150$. The auction has six bidders. Bidders 1, 2 and 3 each want 1 unit of A. Bidders 4, 5 and 6 each want 1 unit of B. As shown in the top panel of Table 1, the bidders have private values for each commodity distributed around the seller's value. The seller uses a discriminatory auction and sets a reserve price of 25 for both commodities. Bids are constrained to be integers, ties are decided at random and the marginal bid may be accepted partially. In this complete information example, every auction parameter (e.g. values, reference prices) is common knowledge.

The second panel of Table 1 shows that, despite reserve prices, a *combined auction* (i.e. highest bids win regardless of commodity) exhibits severe adverse selection (i.e. the seller's is profit is highly negative) because the lack of effective cross-commodity competition enables bidders to submit low bids at the expense of the seller.

In the accurate reference price auction, the seller normalizes the bids submitted with reference prices equal to his own values V_A and V_B. The highest normalized bids are accepted first until the minimum revenue target is reached. If there are bids left that look favorable to the seller (i.e. with a normalized bid strictly greater than 1), then the seller accepts those bids until he reaches his maximum revenue target. As indicated in the third panel from the top in Table 1, the seller's profit is now positive and much larger than in the combined auction. The accurate reference price auction is also efficient in the sense that the commodities are allocated to the three bidders who value them the most in relative terms, i.e. with the highest values relative to their mean. Finally, the accurate reference price auction generates gains from trade with a total surplus of 14.

Next, consider the case where the seller does not know his own values before the auction. Instead, assume that the seller cannot distinguish A from B and initially believes they are both worth 47 (instead of 55 and 39). The seller then sets the reference price for each commodity equal to this value estimate. In equilibrium, bidders 1, 2 and 3 take advantage of the undervalued reference price on commodity A by bidding just above the inaccurate reference price which allows them to clinch one unit each. The seller profit becomes negative, the allocation is inefficient, and total surplus is reduced to 0 (Table 1 fourth panel from the top). This example illustrates how auction performance deteriorates when reference prices are noisy. However, observe that the average submitted bids for each commodity (48 for A; 39 for B) should have provided the seller an indication that A is more valuable than B, and that setting the same reference price for both commodities was a mistake. This is the core idea behind the endogenous reference price auction.

In the endogenous reference price auction, the seller still initially believes the two commodities have an equal value of 47. The seller, however, now revises his

Table 1 Illustrative example

Revenue Targets: *Min*=125, *Max*=150. Reserve Price: 25 for A and B.

		Commodity A			Commodity B			Seller's Profit	Bidders' Profit	Efficiency Measure[†]	Total Surplus
		Bidder			Bidder						
		1	2	3	4	5	6				
	Seller's Value		55			39					
	Bidder's Value	59	55	51	47	41	29				
	Bidder's Relative Value	1.07	1.00	0.93	1.21	1.05	0.74				
Combined Auction	Bid	25✓	25✓	25✓	25✓	25✓	25✓	-132	+132	1.00	0
	Bidder Profit	34	30	26	22	16	4				
Accurate Reference Prices	Bid	56✓	55	51	40✓	40✓	29	+ 3	+11	1.11	+ 14
	Reference Price		55			39					
	Normalized Bid	1.02	1.00	0.93	1.03	1.03	0.74				
	Bidder Profit	3	0	0	7	1	0				
Inaccurate Reference Prices	Bid	48✓	48✓	48✓	47	41	29	- 21	+21	1.00	0
	Reference Price				47						
	Normalized Bid	1.02	1.02	1.02	1.00	0.87	0.62				
	Bidder Profit	11	7	3	0	0	0				
Endogenous Reference Prices	Bid	52✓	52✓	51	46✓	41	29	+1	+11	1.09	+12
	Value Estimate				47						
	Median Bid		52			41					
	Reference Price		49.5			44					
	Normalized Bid	1.05	1.05	1.03	1.05	0. 93	0.66				
	Bidder Profit	7	3	0	1	0	0				

[†] The "Efficiency measure" is the average relative value of the bids accepted. ✓ indicates a bid accepted by the auctioneer.

initial estimate by setting the endogenous reference price for a commodity mid-way between the initial estimate of 47 and the median of the bids submitted for that commodity. As a result, the reference price for the low-quality commodity B is revised downward, which makes the bid of bidder 4 competitive and forces bidders 1 and 2 to bid more aggressively. The seller's profit, the efficiency measure and total surplus all improved substantially compared to the case of inaccurate reference prices, and these measures are close to the best case scenario of accurate reference prices.

To summarize, this example illustrates how noisy seller's value estimates can lead to problems that can be partially corrected with a trivial endogenous process. This result is not contingent on the auction features chosen for this example. In particular, it is easy to verify that the bids in Table 1 remain in equilibrium when the seller uses a uniform pricing rule.[15] This shows that the allocation mechanism, *not* the pricing mechanism, can be of first order importance for auction performance. Similarly, the performance ranking across allocation mechanisms remains unchanged under truthful bidding. This shows that the problems stemming from noisy seller's value estimates we address in this paper *cannot* be solved simply by inducing aggressive/truthful bidding.

---

[15] AHP explain how commodity specific uniform prices are set in a reference price auction: The "market clearing normalized bid" (i.e. the lowest accepted or the highest rejected normalized bid) is first calculated. This common "market clearing normalized bid" multiplied by the reference price of a commodity then determines the "market clearing price," i.e. the commodity specific uniform price per unit. The bids in Table 1 remain in equilibrium whether the market clearing normalized bid is set by the lowest accepted or the highest rejected normalized bid.

5 The experiment

5.1 The treatments

We conduct a between-subjects experiment in which fixed groups of four subjects compete in a series of twelve discriminatory reference price auctions in which bidders have affiliated private values. In each auction, four subjects can bid on four commodities, labeled A, B, C, and D. Following, the classic design of Kagel et al. (1987), the value of a unit of commodity $j \in \{A, \dots, D\}$ to the seller, V_j, is drawn from a uniform distribution with range [\$20,\$80]. The value of any unit of commodity j to bidder i, v_{ij}, is then drawn independently from a uniform distribution with range of ±\$10 around V_j. A bidder can submit up to 6 bids on each commodity, where a bid is a price for a unit of the commodity. Bids are integer-valued between a minimum (reserve price) of \$10 and the bidder's value (since bidding above value is strictly dominated).[16] The seller's minimum and maximum revenue targets are $Min=$ \$1200 and $Max=$ \$2000.

The experiment consists of two sets of treatments under partial and full disclosure. We start with describing the three treatments under partial disclosure. In the "Known" treatment, which serves as the best-case scenario and baseline treatment, the seller knows his values V_j and sets reference prices accurately ($R_j = V_j$). In the "Noisy" treatment, the seller does not know V_j at the time of the auction. Instead, the seller receives a noisy signal $\widehat{V}_j$ for each commodity drawn from a uniform distribution with range of ± \$15 around V_j. The reference price for commodity j is then set to be this noisy estimate ($R_j = \widehat{V}_j$). In the "Endogenous" treatment, the seller receives the same noisy signals $\widehat{V}_j$ as in the Noisy treatment, but the endogenous reference price for commodity j is now $R_j = 0.5\widehat{V}_j + 0.5\bar{b}_j$, where $\bar{b}_j$ is the median of the bids submitted for commodity j.

The same information is provided in all three treatments. In particular, the experimental instructions are identical, subjects know the seller's revenue targets Min and Max, the distribution of the seller's value V_j and the distribution of the bidder's values v_{ij}. Finally, at the beginning of each auction, subjects are told their own value v_{ij} for each commodity.

The same information is also withheld in all three treatments. In particular, subjects at the time of bidding do not know the other bidders' values, the seller's value V_j, the seller's noisy signal $\widehat{V}_j$, the distribution of $\widehat{V}_j$, the seller's reference price R_j and the way R_j is calculated. Thus, subjects do not know that the seller's reference prices can be accurate, noisy or endogenous. Subjects are simply told that the seller uses the "best available information" to determine R_j.

After each auction, the same feedback is provided in all three treatments. The reference prices, normalized bids, and the acceptance status for every bid submitted at the auction are announced. The seller's values V_j and the other bidders' values are not revealed, either before or after the auction. Thus, subjects cannot infer how the

---

[16] There is considerable evidence that restricting bidding above value in discriminatory private values auction experiments does not affect behavior meaningfully (e.g. Cox, Smith and Walker, 1985; Kagel et al., 1987; Kagel, 1995).

reference prices were determined, or whether the seller had accurate or noisy value estimates.

Because the same partial information is provided in the three treatments, subjects are unaware of any treatment differences. This is an important feature of the design which, as discussed earlier, is consistent with many practical auctions that only partially disclose the allocation process (see Sect. 3). Nevertheless, we recognize that withholding relevant information about the game from subjects raises valid deception concerns.[17] To address these concerns, and to test the robustness of our approach, we conduct two additional treatments under full disclosure. The "Explained Noisy" treatment is identical to the Noisy treatment, except that subjects are told that the seller's reference price for a commodity j is drawn from a uniform distribution with range of $\pm$\$15 around the seller's value V_j. The "Explained endogenous" treatment is identical to the Endogenous treatment, except that subjects know how each endogenous reference price is set midway between the seller's own noisy estimate and the median of the submitted bids. Because the partial disclosure treatments are more likely to be implemented in practice, we focus on these treatments in the next sections, and we discuss the results of the full disclosure treatments in Sect. 9.

Each of the five treatments consists of 6 independent sessions, each with 12 auctions and 4 bidders. The same sequence of random draws is used so that each of the 12 auctions is directly comparable across sessions and treatments. Each session lasted for about an hour. A total of 120 subjects were recruited for the five treatments. Subjects were paid a show-up fee of \$10, with subsequent earnings determined by a fractional payout (0.08) of the subject's earnings for all 12 auctions. Earnings averaged about \$35, plus the show-up fee.

5.2 Theoretical predictions under partial disclosure

Because subjects do not know how the seller sets the reference prices in the partial disclosure treatments, they need to form beliefs. We assume a common belief structure across bidders and treatments under which bidder i expects the reference price for object j to be the expected seller's value V_j conditional on the bidder's own signal v_{ij}: $E_i[R_j] = E[V_j|v_{ij}]$. This assumption is consistent with the fact that in the best-case scenario, the seller sets the reference prices equal to his own values: $R_j = V_j$. Nevertheless, we acknowledge that the theoretical predictions in this section are only relevant under this relatively strong assumption.

In principle, the game played in the experiment can be solved numerically because the number of values and bids are finite. The dimension of the strategy space, however, is extremely large. With 4 objects and 6 bids per object, a bid function is a correspondence

---

[17] There is no consensus about what constitute deception in experimental economics. In particular, Krawczyk (2019) as well as Charness, Samek and van de Ven (2022) report that practitioners unambiguously agree that lying to subjects is deceptive, but they consider withholding relevant information a gray area.

between a 4-tuple of values and a 24-tuple of bids, with each component taking integer values between 10 and 90. So the number of possible actions exceeds 10^{35}. To reduce the bid function to a manageable size, we focus on non-decreasing strategies, we constrain values and bids to be even numbers, and we allow each bidder to submit only two bids per object, the first for the first 3 units and the second for the last 3 units. Then, we can use simulations to determine the risk neutral Bayesian Nash equilibrium.[18]

The experimental treatments can be compared according to three performance measures: The seller's profit, total surplus, and allocative efficiency. The seller's profit is the difference between the sales revenue and the value to the seller, $b_{ij} - V_j$, over all the bids b_{ij} accepted at an auction. Total surplus is the sum of the seller's and bidders' profits, $v_{ij} - V_j$, over the bids b_{ij} accepted at an auction. To define allocative efficiency we need to account for differences in values across commodities as well as differences in the number of units accepted across auctions. To do so, we focus on the bidder's relative value (v_{ij}/V_j) averaged over the bids accepted at an auction as in Sect. 4.[19] We consider two allocative efficiency benchmarks. Under the lower benchmark, bids are accepted at random. Because bidders' values are equal to V_j on average, the average relative value of accepted bids is equal to 1 under the lower benchmark. Under the upper benchmark, the bidders bid truthfully ($b_{ij} = v_{ij}$) and bids are normalized with accurate reference prices ($R_j = V_j$). The normalized bids (b_{ij}/R_j) are equal to the bidders' relative values in that case, and bids are accepted in decreasing order of relative values. In other words, under the upper benchmark, the commodities are systematically allocated to the bidders who value them the most in relative terms. In the remainder, the average relative value of the bids accepted at an auction is normalized into a generalized allocative efficiency measure between the lower and upper benchmarks. Thus, an efficiency measure of 0% means that the bids were accepted at random, while a measure of 100% implies full efficiency.

The three performance measures under equilibrium behavior in the known, noisy and endogenous treatments can be found in the top panel of Table 2. We also add as a reference the performance measures obtained under truthful bidding in the second panel. Observe first that the ranking of treatments with respect to each performance measure is identical under equilibrium and truthful bidding, and it is consistent with the ranking we obtained in the complete information example in Sect. 4. Thus, we are now able to make three predictions, P_1, P_2 and P_3, to be tested in the experiment: The seller's profit (prediction P_1), allocative efficiency (prediction P_2), and total surplus (prediction P_3) are (i) highest in the Known treatment, (ii) lowest in the Noisy treatment and (iii) substantially higher in the Endogenous treatment when noisy reference prices are revised.

---

[18] The intuition behind the algorithm used to find equilibrium bid functions is that one can exploit the extensive form of the discrete game and the monotonicity of the strategies to test in sequence (from least to most aggressive) whether a bidding strategy is in equilibrium for all possible value combinations.

[19] Although the two concepts are related, it is easy to see that the allocation that maximizes allocative efficiency (subject to the seller's budget constraints) is different from the allocation that maximizes total surplus.

6 Results from the partial disclosure treatments

6.1 Bid function

We start by comparing in Fig. 1 the bids submitted in the experiment (cross markers) with the equilibrium bid function calculated in the Sect. 5.2 (solid lines). To simplify, we combine all three treatments when plotting the experimental data, consistent with the fact that subjects were unaware of any treatment difference. Figure 1 then displays the average bids submitted on the 4 objects, ordered from highest to lowest bidder's value. For instance, the lowest marker in the top-left panel of Fig. 1, located at $x = 49$ and $y = 44.3$, indicates that the average bid (across units, bidders, auctions, sessions and treatments) submitted in the experiment is 44.3 when the highest of the four values a bidder receives is 49. Likewise, the highest marker in the bottom-right panel of Fig. 1, located at $x = 47$ and $y = 42.4$, indicates that the average bid submitted in the experiment is 42.4 when the lowest of the four values a bidder receives is 47. Similarly, we plot in each panel of Fig. 1 the average equilibrium bid functions for the corresponding object.[20]

Figure 1 shows that, despite the complex strategic environment, bidding behavior in the experiment is reasonably close to the equilibrium bid function.[21] Figure 1 also suggests that, consistent with the experimental literature on single-unit auctions (e.g. Cox et al., 1983), subjects tend to overbid (i.e. submit bids above the risk neutral equilibrium) on the object with the highest value. In contrast, subjects tend to underbid on the object with the lowest value.

6.2 Seller's profit

Figure 2 shows the seller's profit in each auction and each treatment averaged across the 6 independent sessions. Recall that, comparing treatments auction-by-auction is meaningful because the same random draws are used across treatments. Treatment effects (per-auction averages) can be seen in the third panel from the top in Table 2 and by comparing the bars on the right side of Fig. 2. The statistical significance of treatment effects throughout the paper is determined by comparing the distributions of session averages across treatments using a 2-sided permutation test (a non-parametric, exact test) of differences in means (see Good, 2000).

---

[20] Formally, if we order a bidder's values for each object such that $v_{(1)} \geq v_{(2)} \geq v_{(3)} \geq v_{(4)}$, then the equilibrium bid function calculated in Sect. 5.2 is an eight dimensional vector where each component $B^*_{j,k}\left(v_{(1)}, v_{(2)}, v_{(3)}, v_{(4)}\right)$, with $j = 1, 2, 3, 4$ and $k = 1, 2$, represents the bid submitted for the first (when $k = 1$) or last (when $k = 2$) three units of the object with the jth highest bidder's value. We plot in each panel of Fig. 1 $B^*_j\left(v_{(j)}\right)$, the average equilibrium bid submitted on object $j = 1, 2, 3, 4$ *when* $v_{(j)}$ is the jth highest value. For instance, $B^*_1(90)$ is the average of $B^*_{1,1}\left(90, v_{(2)}, v_{(3)}, v_{(4)}\right)$ and $B^*_{1,2}\left(90, v_{(2)}, v_{(3)}, v_{(4)}\right)$ across all possible values of $v_{(2)}, v_{(3)}$ and $v_{(4)}$ that satisfy $90 \geq v_{(2)} \geq v_{(3)} \geq v_{(4)}$. Similarly, $B^*_3(50)$ is the average of $B^*_{3,1}\left(v_{(1)}, v_{(2)}, 50, v_{(4)}\right)$ and $B^*_{3,2}\left(v_{(1)}, v_{(2)}, 50, v_{(4)}\right)$ across all possible values of $v_{(1)}, v_{(2)}, v_{(4)}$ that satisfy $v_{(1)} \geq v_{(2)} \geq 50 \geq v_{(4)}$.

[21] However, the hypothesis that subjects bid according to the equilibrium bid function is easily rejected.
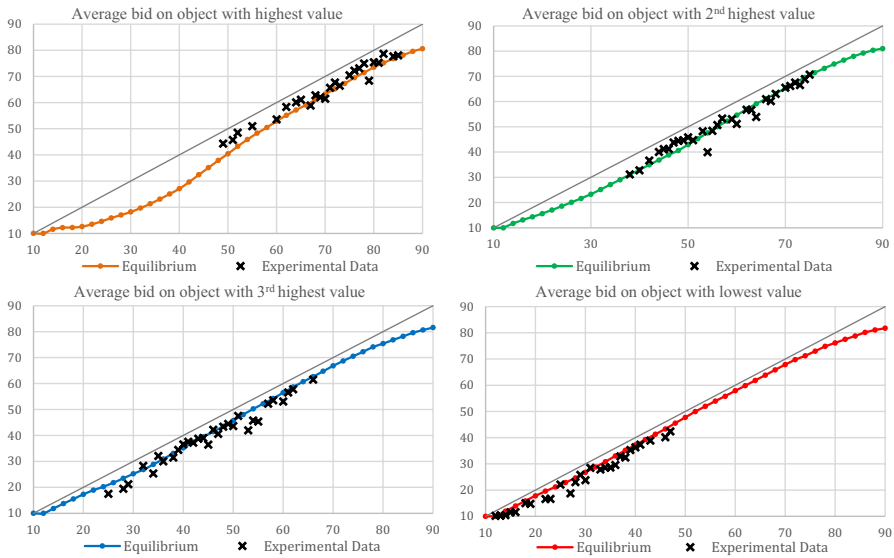
⛄ Springer

Fig. 1 Equilibrium bid function and bidding behavior in the experiment. The Figure plots the average bids submitted in all three treatments of the experiment (cross markers) along with the average equilibrium bid function (solid lines). At an auction, a bidder receives 4 values, one for each for each object. The values can be ordered from highest to lowest. The top left panel represents the average bid (i.e. averaged across units, bidders, auctions, sessions and treatments) submitted on the object with the highest value. The top right panel represents the average bid submitted on the object with the second highest value. The bottom left panel represents the average bid submitted on the object with the third highest value. Finally, the bottom right panel represents the average bid submitted on the object with the lowest value

The seller's profit in the Known treatment is $60.7 per auction on average, −$47.5 in the Noisy treatment and $16.2 in the Endogenous treatment, as indicated in Table 2. The difference in seller's profit across any pair of treatments is significant ($p < 0.05$). Table 2 also shows that the average seller's profit in each treatment is lower than predicted by theory, but the ranking of treatments with respect to the seller's profit is consistent with equilibrium behavior. Figure 2 indicates that the treatment effects are almost systematic: The seller's profit in the Known treatment is positive in 10 of the 12 auctions, and it is higher than in the other two treatments in all but one auction. Further, the seller's profit is always higher in the Endogenous treatment than in the Noisy treatment, and it is significantly higher (at the 5% level) in 9 of the 12 auctions. Note also that the gains from endogenizing reference prices are substantial for the seller. Indeed, when compared with the seller's profit in the Known treatment, the total losses caused by noisy value estimates are reduced from $108.2 in the Noisy treatment ($60.7 – (–$47.5)) to $44.5 in the Endogenous treatment ($60.7 – $16.2). In other words, the seller's losses due to noisy value estimates are reduced by nearly 60% when reference prices are endogenized. This leads to our first result.
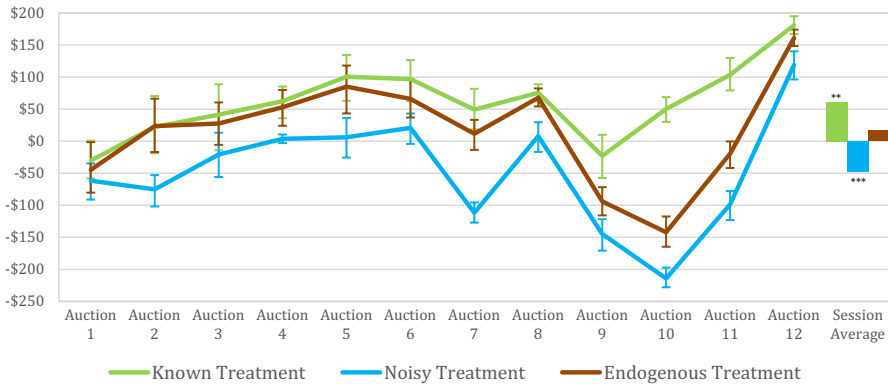
🍁 Springer

Fig. 2 Seller's profit. The points on a line represents the average of the 6 sessions conducted for the corresponding treatment in each of the 12 auctions. Recall that the points in each auction are directly comparable across treatments. The error bars around each point represent the 95% confidence interval. Bars on the right of the figure are per-auction treatment averages. The stars above the green (respectively, blue) bar represent the outcome of a 2-sided permutation test of equal means between the Endogenous treatment and the Know (respectively, Noisy) Treatment. The superscripts ***, **, and * indicate that the null hypothesis of equal means is rejected at the 1, 5, and 10% significance levels

> **Result 1**: Consistent with prediction P1, the seller's profit is significantly higher when the seller knows how much he values each commodity. When the seller has noisy value estimates, a simple endogenization process increases the seller's profit significantly.

6.3 Allocative efficiency

Figure 3 shows the allocative efficiency measure in each auction and each treatment. We can see that the reference price auction consistently produces high allocative efficiency when reference prices are accurate (85.4% on average). Efficiency, however, drops significantly (to 59.2% on average) when reference prices are noisy and unadjusted ($p < 0.01$). Figure 3 shows that these results are relatively systematic: Allocative efficiency is highest in the Known treatment and lowest in the Noisy treatment in 10 of the 12 auctions. Consistent with the theoretical prediction in Table 2, we find that using a simple endogenization process restores allocative efficiency to a great extent. Figure 3 shows that allocative efficiency in the Endogenous treatment (79.2% on average) is significantly higher than in the Noisy treatment ($p < 0.05$), and only slightly (but not significantly) lower than in the Known treatment. This leads to our second result.

> **Result 2**: Consistent with prediction P2, allocative efficiency is highest when the seller knows his values and lowest when the seller has noisy (unadjusted) value estimates. Allocative efficiency is mostly restored with endogenous reference prices.

🖄 Springer

Table 2 Auction performance measures

		Seller's profit	Allocative efficiency (%)	Total surplus	Reference price Inaccuracy (%)	Bid discount (%)
Prediction under equilibrium behavior	Known	63.2	89.8	202.6	0.0	10.5
	Noisy	−36.5	55.8	121.9	18.1	
	Endogenous	19.4	82.3	166.1	6.5	
Prediction under truthful bidding	Known	123.4	100	222.3	0.0	0.0
	Noisy	47.8	73.6	150.0	18.1	
	Endogenous	104.9	91.0	183.4	5.6	
Experimental results under partial disclosure	Known	60.7**	85.4	186.1*	0.0***	12.0
	Noisy	−47.5***	59.2**	111.7**	17.6**	13.3
	Endogenous	16.2	79.2	157.5	5.8	11.6
Experimental results under full disclosure	Explained noisy	−40.2$^{\times\times\times}$	65.1$^{\times\times}$	116.5$^{\times\times\times}$	13.9$^{\times\times}$	12.2
	Explained endogenous	34.1	80.2	164.3	5.6	10.8

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. The "Explained Noisy" and "Explained Endogenous" treatments are identical to the Noisy and Endogenous treatments except that subjects were told explicitly how the seller sets reference prices. The "Seller's Profit" is the difference between sales revenue and market value of sales to the seller. "Allocative Efficiency" is the average relative value of the bids accepted at an auction normalized between 0% (the bids were accepted at random) and 100% (the commodities are systematically allocated to the bidders who value them the most). "Total Surplus" is sum of the seller's and bidders' profit at an auction. 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 difference between a bidder's value and the bid submitted, expressed as a percentage of value. Treatment effects are evaluated with a 2-sided permutation test of equal means. The superscripts ×× ×, × × ×, and × indicate that the null hypothesis of equal means between the Endogenous treatment and the corresponding treatment is rejected at the 1, 5, and 10% significance levels. The superscripts ***, **, and * indicate that the null hypothesis of equal means between the Endogenous treatment and the Explained Endogenous treatment and the Explained Noisy treatment is rejected at the 1, 5, and 10% significance levels
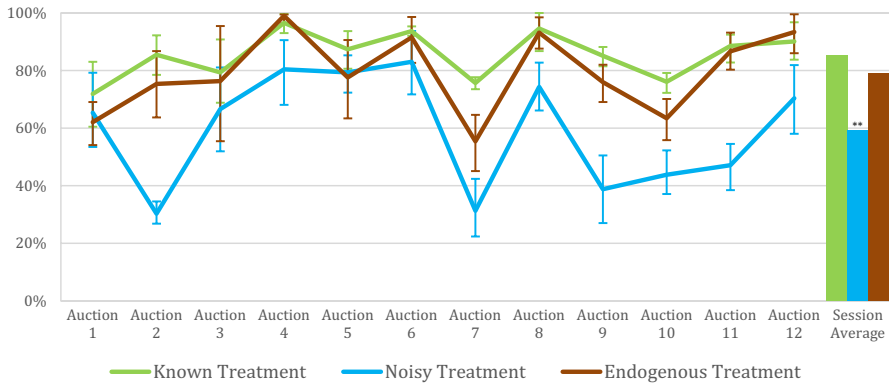
Fig. 3 Allocative efficiency. The points on a line represents the average of the 6 sessions conducted for the corresponding treatment in each of the 12 auctions. Recall that the points in each auction are directly comparable across treatments. The error bars around each point represent the 95% confidence interval. Bars on the right of the figure are per-auction treatment averages. The stars above the green (respectively, blue) bar represent the outcome of a 2-sided permutation test of equal means between the Endogenous treatment and the Know (respectively, Noisy) Treatment. The superscripts ***, **, and * indicate that the null hypothesis of equal means is rejected at the 1, 5, and 10% significance levels

6.4 Total surplus comparison

We can see in Fig. 4 that total surplus is highest when reference prices are accurate (186.1 on average) and lowest when references prices are noisy and unadjusted (111.8 on average). When reference prices are endogenized, total surplus (157.5 on average) is intermediate between the other two treatments: it is significantly higher than in the Noisy treatment ($p < 0.01$), and significantly lower than in the Known treatment ($p < 0.05$). Thus, the treatment ranking with respect to total surplus is consistent with prediction *P3*. This leads to our third result.

> ***Result 3***: Consistent with prediction P3, total surplus is highest with accurate reference prices. Endogenous reference prices produce significantly higher surplus than noisy (unadjusted) reference prices.

To sum up, results 1, 2 and 3 confirm that the reference price auction performs well when the seller knows his own values and set accurate reference prices, but that performance can deteriorate substantially when the seller does not know his values and has to rely on noisy value estimates to set reference prices. This is consistent with the results AHP found in a different context (a uniform price common value auction). Further, consistent with theory, the experimental results show that even a trivial endogenous procedure can help restore most of the auction performance.
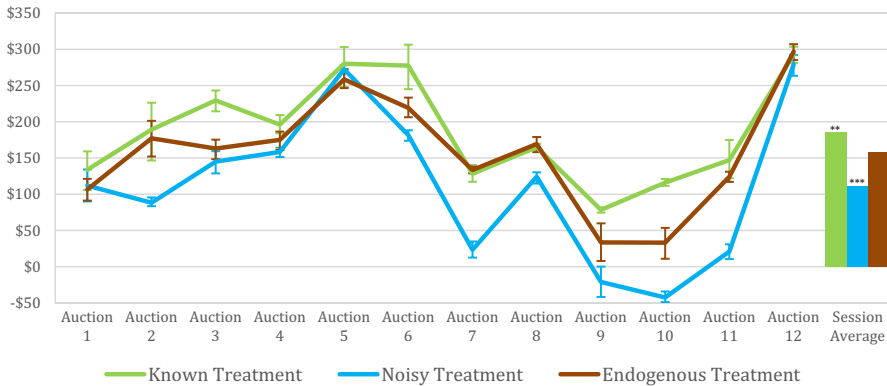
Fig. 4 Total surplus. The points on a line represents the average of the 6 sessions conducted for the corresponding treatment in each of the 12 auctions. Recall that the points in each auction are directly comparable across treatments. The error bars around each point represent the 95% confidence interval. Bars on the right of the figure are per-auction treatment averages. The stars above the green (respectively, blue) bar represent the outcome of a 2-sided permutation test of equal means between the Endogenous treatment and the Know (respectively, Noisy) Treatment. The superscripts ***, **, and * indicate that the null hypothesis of equal means is rejected at the 1, 5, and 10% significance levels

7 How does endogenization increase performance?

In this section we explore two factors that may explain why auction performance improved with endogenized reference prices: bidding adjustments and a reduction in reference price inaccuracies.

7.1 Bid discounts

Recall that subjects in the experiment are unaware of any treatment. In particular, they do not know how the reference prices are determined and whether or not they are accurate. Thus, in principle, bidding behavior should be identical across treatments. In practice, however, as subjects gain experience, they may learn to behave differently in each treatment. Such bidding adjustments could help explain the difference in auction performance documented in Sect. 6. To test this hypothesis, we define the "bid discount" as the percent difference between the bidder's value and the bid submitted: $(v_{ij} - b_{ij})/v_{ij}$. For instance, a bid discount of 10% indicates that the bidder submits a bid 10% below value. The bid discount therefore captures how aggressive the bidding is (low bid discounts correspond to more aggressive bidding).

Figure 5 shows the average bid discount in each auction and each treatment. Observe first that despite the complex auction environment, bidding behavior stabilizes after only a few auctions in each treatment. Figure 5 also confirms that bidding at the beginning of the experiment is statistically indistinguishable across treatments. After the third auction, however, the average bid discount is systematically highest in the Noisy treatment, although the difference is most often insignificant. Thus, it
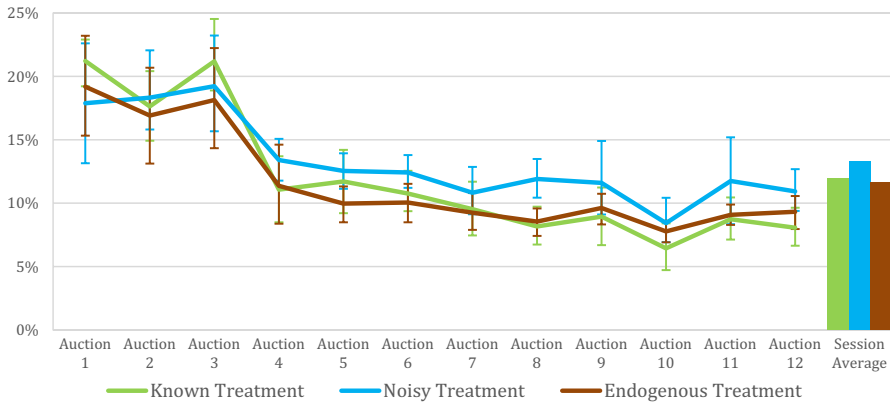
🐾 Springer

Fig. 5 Bid discounts. The points on a line represents the average of the 6 sessions conducted for the corresponding treatment in each of the 12 auctions. Recall that the points in each auction are directly comparable across treatments. The error bars around each point represent the 95% confidence interval. Bars on the right of the figure are per-auction treatment averages. The stars above the green (respectively, blue) bar represent the outcome of a 2-sided permutation test of equal means between the Endogenous treatment and the Know (respectively, Noisy) Treatment. The superscripts ***, **, and * indicate that the null hypothesis of equal means is rejected at the 1, 5, and 10% significance levels

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. It is also interesting to note that subjects are initially less aggressive than predicted under equilibrium behavior (i.e. the average bid discount in the first three auctions of every treatment exceeds the 10.5% equilibrium bid discount in the last column of Table 2). With experience however, subjects in all treatments except the Noisy treatment learn to bid more aggressively than theory predicts (i.e. the average bid discount in the last three auctions becomes lower than 10.5%). We can now state our fourth result.

> **Result 4**: Bidding behavior was on average less aggressive in the Noisy treatment, but statistically indistinguishable in the Known and Endogenous treatments.

7.2 Does endogenizing reference prices reduce inaccuracies?

To measure the inaccuracy of a reference price, we calculate the absolute percentage difference between the seller's value and the reference price actually used to normalize the bid of the corresponding commodity: $\left| R_j - V_j \right| / V_j$. We can see in Table 2 that this reference price inaccuracy measure is nearly three times smaller in the Endogenous treatment than in the Noisy treatment (5.8 versus 17.6% on average). This treatment effect is highly significant ($p < 0.01$), and it is consistent with the prediction under equilibrium behavior in Table 2. This leads to our fifth result.

Result 5: The reference prices in the Endogenous treatment are significantly more accurate than the unadjusted reference prices in the Noisy treatment, consistent with theory.

To sum up, the combination of results 4 and 5 confirms that, as suggested by theory, most of the increase in auction performance between the Noisy and Endogenous treatments comes from more accurate reference prices, not from adjustments in bidding behavior.

8 Alternative endogenization processes

The experimental results show that even a simple endogenization of the reference prices can improve auction performance. In this section, we use equilibrium behavior to explore the extent to which performance can be increased further by using a more sophisticated endogenization process.

8.1 Weighted average

We begin by considering a more general process in which the endogenous reference price for commodity j is a *weighted* average of the seller's noisy value estimate $\widehat{V}_j$ and the median bid $\overline{b}_j : R_j\left(\overline{b}_j\right) = w * \overline{b}_j + (1 - w) * \widehat{V}_j$, with $0 \leq w \leq 1$. The seller ignores the bids submitted when $w = 0$, and only the median bid is taken into consideration when $w = 1$. In the Noisy and Endogenous treatments, the weight w is set to 0 and ½, respectively. Figure 6 shows how the three performance measures change with w under equilibrium behavior (solid lines with circle markers). As upper and lower benchmarks, we also plot the predicted performance measures in the Known and Noisy treatments (upper and lower horizontal lines, respectively). Ignore for the moment the dotted and dashed lines which will be discussed in the next section.

Relatively similar patterns emerge for the three performance measures in Fig. 6: Performance initially increases when some weight is given to the median bid, but it ultimately decreases when the seller gives insufficient weight to his noisy signal. For the specific set of parameters used in the experiment (i.e. number of bidders, values and signals' distributions), a weight around 0.7 seems to yield the highest performance measures. In particular, $w = 2/3$ would yield the seller a profit of \$25.0 per auction on average in equilibrium. This is higher than the predicted seller profit of \$19.4 in the Endogenous treatment (indicated in Fig. 6 by a full dot at $w = 0.5$), but this remains substantially lower than the prediction of \$63.2 in the Known treatment (upper horizontal line). Thus, using a weight around 0.7 further mitigates the impact of noisy value estimates, but it does not solve the problem entirely.
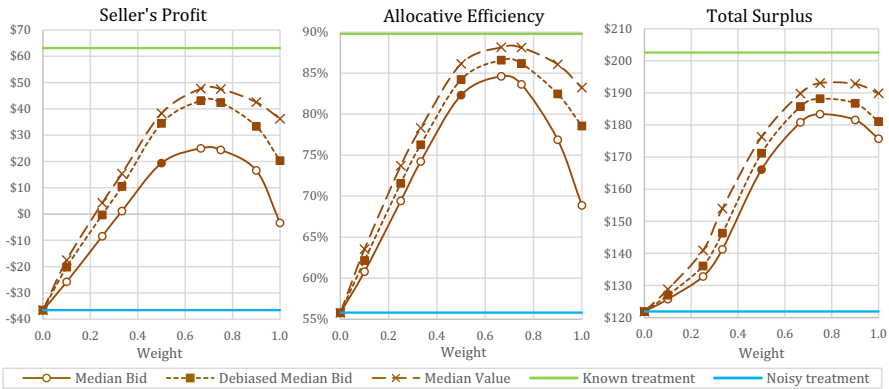
🖄 Springer

Fig. 6 Alternative endogenous process. The connected hollow dots show predicted performance measures 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 ranges from 0 on the left side of the horizontal axis to 1 on the right side. The full dots at $w=\frac{1}{2}$ represent the predicted performance measures in the Endogenous treatment. The lower and upper flat lines indicate the predicted performance measures in the Noisy and Known treatments, respectively. The dotted lines connecting the square markers plot predicted performance measures when the median bid is debiased before using the weighted average formula. The dashed lines connecting the cross markers show predicted performance measures when reference prices are based on the median *bidder value* (instead of the median bid)

8.2 Weighted average with bias correction

Although simple and easily implementable, our approach to endogenize reference prices based on median bids is in no way optimal. In particular, Table 2 shows that bidders discount their bids by 10.5% on average under equilibrium behavior. Thus, the median bids may be reasonable and robust estimates of the seller's values, but they are nevertheless biased downward and so are the endogenous reference prices.[22] Can the seller reduce this bias?

In principle, a correction is possible if the bids are in equilibrium and the auctioneer is able to invert the equilibrium bid function to recover the bidders' values from the bids they submit. Given the complexity of the environment, such an approach may not be feasible in practice. Alternatively, the auctioneer could learn to correct the bias in the median bid from observing the outcome of past auctions. For example, if the seller's actual value tend to be 15% above the median bid in past auctions, then the seller could apply a 15% "debiasing factor" to the median bid in future auctions.

---

[22] Klemperer (2010) argues that truthful bidding is a reasonable assumption in a (highest rejected) multi-unit uniform price auction because no single bid is likely to set the auction clearing price. Thus, the median bid is an unbiased estimate of the seller's value in this case and no bias correction would be necessary.

To gauge the effectiveness of this simple debiasing method, we first calculate for each auction $t = 1, \ldots, 12$ the debiasing factor δ_t equal to the average ratio of the seller's value to the median bid across all commodities sold at previous auctions: $\delta_t = \frac{1}{4*(t-1)} \sum_{\tau=1}^{t-1} \sum_{j=1}^{4} \left(V_{\tau j} / \overline{b}_{\tau j} \right)$. 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 to the weighted average of the debiased median bid and the seller's noisy value estimate: $R_{tj}\left(\overline{b}_{tj}, \delta_t \right) = w * \left(\delta_t * \overline{b}_{tj} \right) + (1 - w) * \widehat{V}_{tj}$. For the first auction, we set $\delta_1 = 1$ because the seller has no information to debias the median bid.

Figure 6 shows the predicted performance measures under equilibrium behavior when the weight given to the "debiased" median bid varies from 0 to 1 (dotted lines with square markers). The figure indicates that endogenizing the reference prices using the debiased median bids increases auction performance for every value of the weight w. Using the debiased median bids and a weight around 2/3 yields the highest performance measures. In particular, a weight of $w = 2/3$ yields the seller a profit of \$43.2 per auction on average in equilibrium. This is more than twice as high as the \$19.4 predicted in the Endogenous treatment and 73% higher than the \$25.0 we calculated previously with the same weight but without debiasing the median bids.

This exercise suggests that auction performance can be increased substantially by using data from past auctions to debias the median bid. This result is relevant for the auctions discussed in introduction (e.g. buyback and QE auctions) because they all involve frequently repeated auctions, which gives the auctioneer an opportunity to learn how to make appropriate bias corrections.

8.3 Performance bounds when bidders' values are observed

What would auction performance be under a scenario in which the auctioneer can actually observe the bidders' values? To answer this question, we now assume that the seller revises the reference prices for each commodity by taking a weighted average of his noisy value estimate and the median of the *bidders' values* (instead of the median bid).

Figure 6 shows the predicted performance measures when the weight given to the median of the bidders' values varies from 0 to 1 (dashed lines with cross markers). As expected, performance is systematically higher than in the other two exercises. Further, observe that the seller should give most of the weight to the median value, but he should not ignore his own noisy estimate (which would correspond to the case $w = 1$). More importantly, note that the three performance measures produced by the bias correction method described in the previous section are very close to the bounds that would be obtained if the seller knew the bidder's values. In particular, the average seller's profit of \$43.2 obtained for $w = 2/3$ with the bias correction method is close to the \$47.7 upper bound obtained using the median values. Thus, when the seller does not know his values, alternative methods to endogenize the reference prices are unlikely to provide major additional benefits compared to the simple bias correction method.

🖄 Springer

9 Endogenous reference price auctions under full disclosure

9.1 Treatments under full disclosure

The results in Sect. 7 suggest that a simple endogenous process can enhance auction performance significantly when the seller does not disclose how reference prices are set. To test the robustness of this result, we now turn to the results of the two treatments under full disclosure which are summarized in Table 2 and Figs. 7 to 14 in Supplementary material.

We find no evidence that telling subjects how the seller sets the reference prices introduces a treatment effect. As shown in Figs. 7 to 10 in Supplementary material, the seller's profit, allocative efficiency, total surplus and bidding discounts are not significantly different in the Explained Noisy and Explained Endogenous treatments than in the Noisy and Endogenous treatments, respectively. Hence, performance does not seem to suffer when subjects know the endogenization process. In fact, a comparison of the last two rows of Table 2 shows that the three performance measures are actually slightly higher (although not significantly) in the full disclosure treatments. More importantly, Table 2 confirms our main result: Under full disclosure, the Endogenous treatment also significantly outperforms the Noisy treatment with respect to each of the three performance measures (see also Figs. 11 to 14 in Supplementary material).

9.2 Avoiding possible bidding manipulation

When the endogenization process is known, a bidder may have an incentive to submit "phantom bids" (i.e. bids at a very low price) with the sole objective of lowering a commodity's median bid. By doing so, the bidder may be able to lower the endogenous reference price and make another one of his bids look more attractive to the seller. The experiment we conducted provides no evidence of such bidding manipulation. In particular, only 0.2% of the bids submitted in the Explained Endogenous treatment correspond to a bid discount that exceeds 35%. Thus, bidders did not try to influence the median bid by submitting some bids substantially below value.

Nevertheless, the possibility of phantom bids exists when the endogenous process is common knowledge. So, how can the seller avoid possible bidding manipulation in practice? One simple way to do so is by not disclosing explicitly the allocation mechanism, an approach often taken by auctioneers in practice (see Sect. 3).[23] Another measure the seller could take to eliminate incentives for submitting phantom bids is to modify the endogenization process. 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

---

[23] The approach, however, may not be fully supported by theory, as general predictions are hard to obtain when the allocation mechanism is not common knowledge.

bids are not taken into consideration.[24] Simulations under truthful bidding indicate that this approach produces slightly higher but substantially more volatile performance measures. For instance, the average and standard deviation of the seller's profit are \$101.6 and \$90.2 under this approach, compared to \$104.9 and \$67.3 in the Endogenous treatment. This result reflects the fact that using only the highest bid submitted by each bidder for a commodity yields an estimate of the seller's value that is slightly less biased (because it relies on the bids closest to each bidder's value), but more variable (because it is calculated with fewer data points). This exercise illustrates how the endogenization process can be adjusted to remove incentives for bidding manipulations without reducing auction performance.[25]

10 Conclusion

An auctioneer using a multi-object, multi-unit auction has to compare bids across different commodities. If the auctioneer knows his own value for each commodity, then the "reference price auction" (Armantier et al., 2013) enables the auctioneer to accept the bids with the best relative values. This paper focuses on the case of an auctioneer with noisy value estimates. We propose a generalization of the reference price auction in which the auctioneer reduces value inaccuracies by setting reference prices endogenously using information extracted from the bids submitted. To illustrate how reference prices may be endogenized, we consider a simple process under which the reference price for a commodity is set midway between the seller's initial noisy value estimate and the median of the bids submitted for that commodity. Consistent with theoretical predictions, the experimental results indicate that even such a simple endogenous process significantly increases three measures of auction performance: the seller's profit, allocative efficiency and total surplus. We also find that using data from past auctions to set reference prices can help the auctioneer raise auction performance to a level that is near the upper bound for any endogenous method. Further, adjustments to the endogenous process or keeping it undisclosed prevent bidding manipulation.

These results have practical implications because many important auctions are both multi-object, multi-unit and involve a budget-constrained auctioneer who may have imperfect value estimates. Among the examples presented in introduction, consider first the Treasury buyback auctions. The General Accounting Office (GAO) has endorsed buyback auctions as a regular and effective debt management tool,

---

[24] There are other ways the seller could modify the endogenization process to avoid manipulations. For instance, the seller could rely only on the median of the *winning bids*. This approach is more complex (the auctioneer has to solve a fixed-point problem to determine the winning bids) but it can be addressed with standard numerical techniques.

[25] Collusion is also a legitimate concern for any auction mechanism. Although not immune to the problem, the endogenous reference price auction discussed in this paper is less susceptible to collusion. Indeed, as explained in AHP, one of the key benefits of a reference price auction is to promote competition by creating thick markets in which more bidders can compete across commodities, thereby making collusive behavior more difficult to sustain.

both in times of surplus and deficit.[26] The GAO also recommends using a multi-object, multi-unit format, and emphasizes the importance of lowering the Treasury's purchasing cost by accepting the most competitive bids. The GAO, however, does not specify how the Treasury should "cherry pick" among the bids submitted to achieve this objective. Further, because the Treasury is not an active trader on the secondary market, it may not be able to assess precisely the relative values of the bids submitted for different bonds. In this context, the endogenous reference price auction provides a simple yet effective solution to this allocation problem. Consider now the massive QE auctions that are increasingly conducted around the world, in particular since the start of the Covid-19 pandemic. Because QE programs are launched in times of crisis when markets tend to be impaired, these auctions present a unique challenge: How can the auctioneer compare bids across different securities when one cannot rely on market prices to assess their fundamental values?[27] Again, the endogenous reference price auction would enable the auctioneer to address this challenge and improve performance at QE auctions.

We believe that the contribution of this paper is not limited to the manner in which we proposed to endogenize reference prices. More generally, it is in the recognition that the auctioneer can exploit information contained in the bids submitted to improve auction performance. This is in contrast with the literature which has focused on the payment mechanism (e.g. discriminatory vs. uniform price) or on trying to promote aggressive truthful bidding. These approaches, however, are of little help to address the problems stemming from noisy seller's value estimates we consider in this paper. Instead, our results suggest that a simple, cost-free mechanism can help the auctioneer acquire information and improve auction performance.

---

[26] Based on the GAO endorsement, the U.S. Treasury has been conducting small-value buyback auctions since 2014. The IMF and World Bank recently made similar recommendations for emerging countries (Akcadag et al., 2015).

[27] This challenge led Song and Zhu (2018) to conclude "Given that QE operations tend to be large and could become a standard tool in monetary policy, any improvement in the implementation mechanism … could bring substantial efficiency gains." (p. 120). See as well Han et al. (2007) and Breedon (2018), who also argue for better auction designs to lower the cost of Treasury buyback and QE auctions.

Springer

References

Akcadag O., Baudouin R., Silva A. and I. Raghavan (2015) Bond buybacks and exchanges: Background note. Washington D.C., World Bank Group.

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

Armantier, O., & Sporn, J. (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., & Cantillon, B. (2008). Properties of scoring auctions. *RAND Journal of Economics, 39*(1), 69–85.

Back, K., & Zender, J. (2001). Auctions of divisible goods with endogenous supply. *Economics Letters, 73*, 29–34.

Bichler, M., & Goeree, J. (Eds.) (2017). *Handbook of Spectrum Auction Design*. Cambridge University Press.

Bolton, G., & Ockenfels, A. (2012). Behavioral economic engineering. *Journal of Economic Psychology, 33*, 665–676.

Breedon, F. (2018). On the transactions costs of UK quantitative easing. *Journal of Banking and Finance, 88*, 347–356.

Cason, T. (1995). An experimental investigation of the seller incentives in the EPA's emission trading auction. *American Economic Review, 85*(4), 905–922.

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

Charness, C., Samek, A., & van de Ven, J. (2022). What is considered deception in experimental economics? *Experimental Economics, 25*, 385–412.

Chen-Ritzo, C., Harrison, T., Kwasnica, A., & Thomas, D. (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., & Walker, J. (1983). Tests of a heterogeneous bidder's theory of first-price auctions. *Economics Letters, 12*, 207–212.

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

Cox J., Smith V. & J. Walker (1985). Expected revenue in discriminative and uniform price sealed bid auctions. In: Smith, V. (Eds) *Research in Experimental Economics Vol. 3*. JAI Press.

Dyer, D., Kagel, J., & Levin, D. (1989) ".Rsolving uncertainty about the number of bidders in independent private-value auctions: An experimental analysis. *The RAND Journal of Economics, 20*(2), 268–279.

Dyer, D., & Kagel, J. (1996). Bidding in common value auctions: How the commercial construction industry corrects for the Winner's Curse. *Management Science, 42*(10), 1463–1475.

Fréchette G. (2016). Laboratory experiments: Professionals versus students, in Handbook of Experimental Economic Methodology, Oxford University Press, 360–390.

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

Garbade, K. (2015). *Treasury debt management under the rubric of regular and predictable issuance: 1983–2012*, Federal Reserve Bank of New York.

Good, P. (2000). *Permutation Tests*. Springer Series in Statistics.

Han, B., Longstaff, F., & Merril, C. (2007). The U.S. treasury buyback auctions: The cost of retiring illiquid bonds. *The Journal of Finance, 62*(6), 2673–2693.

Hortaçsu, A., & McAdams, D. (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.

Kagel, J., & Levin, D. (2001). Behavior in multi-unit demand auctions: Experiments with uniform price and dynamic Vickrey auctions. *Econometrica, 69*(2), 413–454.

Kagel, J., Harstad, R., & Levin, D. (1987). Information impact and allocation rules in auctions with affiliated private values: A laboratory study. *Econometrica, 55*(6), 1275–1304.

Kagel, J., & Levin, D. (2016). Auctions. A survey of experimental research. In: Kagel, J., Roth, A. (Eds.) *Handbook of Experimental Economics Vol. II*. Princeton University Press.

Kagel, J. (1995). Auctions: A survey of experimental research. In: Kagel, J., Roth, A. (Eds) *The Handbook of Experimental Economics*. Princeton University Press.

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

Krawczyk, M. (2019). What should be regarded as deception in experimental economics? Evidence from a survey of researchers and subjects. *Journal of Behavioral and Experimental Economics, 79*, 110–118.

Kwasnica, A., & Sherstyuk, K. (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.

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

Myerson, R. (1981). "Optimal auction design", *Mathematics*. *Operations Research, 6*, 58–73.

Roth, A. (2002). The Economist as engineer: Game theory, experimentation, and computation as tools for design economics. *Econometrica, 70*, 1341–1378.

Roth, A. (2016). Experiments in market design. In: *Handbook of Experimental Economics Volume 2*. Princeton University Press, pp 290–346.

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

Song, Z., & Zhu, H. (2018). QE auctions of Treasuries bonds. *Journal of Financial Economics, 128*, 103–124.

Varian, H. (2008). How to drink the sub-prime kool-aid. *Economists' Voice,5*(5).