Price Adjustment and Competition in US Strategic Petroleum Reserve Drawdowns

Matthew O'Keefe*

June 12, 2025

Abstract

Sales from the US Strategic Petroleum Reserve (SPR) are conducted using basis auctions, in which winning bidders' final payments are indexed to spot prices. This mechanism can eliminate potential inefficiencies arising from private information about the evolution of oil prices between auction and delivery, benefiting the seller (the US government). Reduced form evidence suggests that bids in SPR sales are consistent with pure private values despite significant uncertainty and a deep resale market. In a stylized model calibrated to the 2022 SPR drawdown, the use of basis auctions generates tens of millions of dollars in savings with limited private information.

Keywords: Strategic petroleum reserve; long-term contracts; contingent payment auctions; private information; oil markets

*Vanderbilt University. Email: matthew.okeefe@vanderbilt.edu. The author gratefully acknowledges insightful comments from Vivek Bhattacharya, Pietro Bonaldi, Rob Porter, Eric Richert, and Charles Z. Zheng.

1 Introduction

In 2022, the United States Department of Energy (DOE) auctioned 180 million barrels of crude oil from the Strategic Petroleum Reserve (SPR). A distinctive feature of the auction mechanism used by DOE is *price adjustment*: the price paid by a winning bidder is linked to the value of an oil price index at delivery, which may occur weeks or months after the conclusion of the auction. This paper investigates the implications of this design choice.

Price adjustment mechanisms are ubiquitous in commodities contracts, power purchase agreements, and many other settings. In the context of long-term contracting, a principle motivation for price adjustment is to mitigate parties' incentives to breach or seek renegotiation as opportunity costs evolve over the life of a contract (Goldberg, 1985). Price adjustment can also improve surplus by mitigating moral hazard problems and preventing the dissipation of rents through costly information acquisition (Goldberg and Erickson, 1987; Joskow, 1988). In the context of auction design, price adjustment generates contingent payments, which can reduce information rents (Hansen, 1985; Riley, 1988) and limit distortions from the winner's curse (Skrzypacz, 2013), among other effects.

This paper focuses on the *competitive* implications of price adjustment for potential buyers of SPR crude. An important feature of SPR auctions is that bidders' valuations depend on future oil prices, which are uncertain. Section 2 presents a simple model that illustrates how linking final payments to an oil price index can eliminate the value of private information about future oil prices, so that bidding strategies and allocations depend only on heterogeneity in bidders' idiosyncratic opportunity costs. By transforming a common values auction with more dispersed information to a private values auction with less dispersed information, price adjustment can increase total surplus and DOE's expected revenue. At the same time, price adjustment greatly reduces the incentive for winning bidders to default, which could otherwise be severe enough to undermine DOE's objectives.

The benefits of price adjustment in the model depend crucially on the similarity between the oil price index and the common component of bidders' valuations. If there are important non-price sources of common values—for example, due to uncertain quality or uncertain transportation costs—or if the oil price index is otherwise imperfect, the particular form of price adjustment used by DOE could deliver few practical benefits. Thus, an assessment of the effectiveness of price adjustment in SPR auctions ultimately turns on the empirical question of whether DOE's basis auctions retain a significant common values element or not.

I address this empirical question by testing whether observed bidding behavior in SPR auctions is consistent with private values. Evidence in favor of private values suggests that DOE's current basis auction format delivers greater expected revenue and improved efficiency relative to a fixed price alternative. In turn, such evidence illustrates a more general benefit of price adjustment in long-term contracting environments characterized by competition between potential counterparties. Conversely, evidence against private values would highlight a practical limitation of price adjustment—namely, the difficulty of constructing suitable price indices—and would indicate scope for improvement of DOE's current mechanism.

To motivate the empirical approach, Section 3 describes the institutional environment in greater detail and introduces bidding data taken from the 2022 SPR drawdown. At its core, an SPR release entails a collection of parallel pay-as-bid (discriminatory) basis auctions for differentiated lots of crude oil stored at different sites on the US Gulf Coast.[1] In practice, the strategic environment is complicated by capacity constraints that limit delivery to particular locations as well as high minimum bid quantities that reflect the considerable fixed costs of coordinating crude oil deliveries. These complications make it difficult to explicitly characterize a more realistic bidding equilibrium than the one introduced in Section 2, precluding the use of structural methods to disentangle private and common values.[2]

Given the complexity of the auction environment, I consider two complementary tests that are valid under relatively weak conditions. First, I consider a reduced form test of common values developed for sovereign debt auctions by Nyborg et al. (2002) and Bjønnes (2001).[3] To implement this test, I exploit the fact that the time elapsed between the auction and delivery is lengthy and varies considerably within the sample. I find that bids are no less aggressive and no steeper when this measure of uncertainty is greater, consistent with private values. Second, I consider a test of independent private values (IPV) developed by Hickman et al. (2021) and previously applied in a multiunit auction environment by Richert (2024b). Specifically, I test for within-auction correlation in bids after controlling for participation and key auction characteristics, finding that IPV cannot be rejected in the data. Given the favorable power properties of the Hickman et al. (2021) test, which I confirm in numerical experiments, this negative result can plausibly be viewed as evidence in favor of IPV and hence as evidence for pure private values. Together, these results suggest that private information plays a limited role in SPR auctions, consistent with the model.

To conclude the paper, I aim to benchmark the aggregate effects of price adjustment

[1] The pay-as-bid and uniform price auctions are the two most commonly used auction formats for selling divisible goods (such as sovereign debt or electricity). In the pay-as-bid auction, winning bidders pay their expressed willingness-to-pay for each inframarginal unit.

[2] Bonaldi et al. (2024) decompose private and common values in sovereign debt auctions by estimating a parametric model of a uniform price auction due to Vives (2011). There are no comparable parametric models for the pay-as-bid auction. Recent empirical work on pay-as-bid auctions in the industrial organization literature has exclusively assumed private values (Hortacsu and McAdams, 2010; Hafner, 2023).

[3] Other applications include Elsinger and Zulehner (2007) and Mariño and Marszalec (2023). A similar approach was considered in Hortaçsu (2002).

during the 2022 SPR drawdown in the context of a calibrated uniform price auction model based on Vives (2011). Unlike the DOE's pay-as-bid auction, the uniform price auction admits a tractable class of parametric linear Bayes Nash equilibria even for complex information structures featuring both private and common values. I calibrate the distribution of future oil prices from crude oil futures and bidders' idiosyncratic opportunity costs from observed bids, leaving private information unspecified.[4] The model yields plausible conditions on private information under which price adjustment could have increased revenue by tens of millions of dollars in an environment similar to the 2022 SPR drawdown.

Related literature The paper makes three main contributions to the literature. First, I contribute to the literature on price adjustment in long-term contracting referenced above by highlighting the value of price adjustment as a tool for shaping competition between potential counterparties. Second, I contribute to the auctions literature by further exploring the effects of contingent payments in environments with both common and private values. The majority of literature on contingent payment auctions falls within the private values paradigm and emphasizes the potential reductions in information rents that are achievable with non-cash auction designs, often contrasting this benefit with potential moral hazard problems if the winner is required to undertake costly ex post actions (DeMarzo et al., 2005) or adverse selection (Che and Kim, 2010; Kong et al., 2022). Bhattacharya et al. (2022) differ in considering an auction environment with pure common values and point out that contingent payments can reduce bidders' need to protect against the winner's curse.[5] In this paper's model of SPR sales, both of these effects are present, as well as a third: price adjustment increases the correlation between bid levels and the private components of bidders' valuations, which (at least in this setting) generates an improvement in efficiency.

Finally, I provide the first game theoretic analysis of the SPR drawdown mechanism. In doing so, this paper complements prior work in macroeconomics analyzing whether SPR releases stabilize oil prices (Newell and Prest, 2017; Killian and Zhou, 2020; Stevens and Zhang, 2021).[6] Earlier work in operations research addresses optimal procurement and drawdown policies for petroleum reserves (Teisberg, 1981; Oren and Wang, 1986).

[4]Observed bids are informative under the assumption that outcomes of DOE's pay-as-bid auctions (most importantly, revenue) would not be significantly different under the uniform price format. In general, the pay-as-bid and uniform price auctions are not revenue equivalent and their relative ranking is an empirical question (see, e..g, Ausubel et al., 2014). Few reliable estimates of revenue differences have been obtained due to the difficulty of solving for equilibria in either format with generic value distributions. In a recent paper circumventing this problem, Richert (2024a) finds that the difference in revenue between discriminatory and uniform price auctions is 0.16% for Turkish treasury auctions.

[5]This observation is also made in Skrzypacz (2013). In finance, it is standard to consider environments in which the seller is privately informed rather than the buyers (Nachman and Noe, 1994).

[6]See also Razek et al. (2023) for macroeconomic analysis of the 2022 SPR releases.

2 Understanding price adjustment in SPR sales

This section illustrates the central insights of the paper using a stylized model motivated by the details of SPR sales, which are further described in Section 3. Sales from the SPR can be described as *basis auctions.* I compare the basis auction format to a generic fixed price auction format. The analysis shows how the basis auction format can improve total surplus and limit the information rents earned by bidders, benefiting DOE. I also discuss ex post incentives, focusing on how the basis auction limits winning bidders' incentive to default.

Suppose the government aims to sell exactly Q barrels of crude oil from the SPR using an auction. For the moment, we assume that all bidders are oil refiners who each demand exactly K barrels of crude oil. There are N such oil refiners, who each differ in their opportunity costs of purchasing SPR crude. Suppose refiner i can always purchase K units of crude on the open market at price $p_i = W + u_i$ where W is an oil price index and u_i is a *cost basis* capturing the difference between i's marginal cost of crude and the index price. To simplify the analysis, let $K = 1$ and $Q = K$, so that a single refinery is awarded exactly K barrels of crude oil with probability one in equilibrium. Because refiners' marginal costs are constant, this setup corresponds to the so-called "flat demands" model of Ausubel (2004), which provides a tractable starting point for analyzing multiunit auctions. With $Q/K = 1$ the pay-as-bid auction is strategically equivalent to a first price auction for a single indivisible good, and can therefore be analyzed using standard methods.

Each oil refiner's cost basis u_i is assumed to be privately known. Furthermore, each bidder privately observes an unbiased estimate w_i of the true but unknown ex post index price W. I assume that w_i is drawn independently across bidders from a distribution F_w having log-concave density f_w with support on a bounded interval $[\underline{w}, \bar{w}]$, while u_i is drawn independently across bidders and from a distribution F_u having log-concave density f_u with support on a bounded interval $[\underline{u}, \bar{u}]$. I assume that w_i is independent of u_i. Importantly, log-concavity of f_u and f_w ensures that bidders can be ordered according to a scalar "type" $s_i = w_i/N + u_i$ (Goeree and Offerman, 2003).

Importantly, delivery of crude oil does not occur until weeks or months after the auction has concluded. During this time, conditions in the broader oil market may change significantly. Because bidders' private signals $w_1, \ldots, w_n$ are distributed independently, the mean signal $\bar{w} = \frac{1}{n}\sum_{i \leq n} w_i$ represents the best available forecast of W at the time of the auction.[7] Hence, $\bar{w}$ can be viewed as the *common value* component of bidders' valuations at the time of the auction, while the realization of W remains unknown until delivery.

[7] $\bar{w}$ is the "best" estimate in the sense that $\bar{w}$ is the lowest-variance unbiased estimator of W that can be constructed from the signals $w_1, \ldots, w_n$. This structure induces a game similar to the "Wallet Game" in Bulow and Klemperer (2002).

A central choice facing DOE is the selection of an auction format. One natural candidate mechanism is a fixed price pay-as-bid auction. Under the assumptions on preferences and quantities stipulated above, bidders submit flat demand curves, and a fixed price pay-as-bid auction corresponds to a first price auction. This simplification makes it straightforward to characterize bidders' strategies in a pure-strategy Bayes Nash equilibrium under a few additional restrictions. In particular, we assume that the winning bidder is obligated to take delivery on a known date, that default is impossible, and that there is no reserve price. Under these assumptions, we obtain the following well-known result.

Proposition 1. *Let $y_i = \max_{j \neq i} s_j$. The N-tuple of strategies $(\beta_1(\cdot), \ldots, \beta_N(\cdot))$, where*

$$\beta_i(x) = E[\bar{w} + u_i | s_i = x, y_i \leq x] - E[x - y_i | s_i = x, y_i \leq x], \tag{1}$$

is an equilibrium of the fixed price pay-as-bid auction.

All proofs are contained in the Appendix.

Equation (1) has a standard interpretation. The first term represents i's interim assessment of his ex post opportunity cost of purchasing SPR crude conditional on winning the auction. This term reflects bidders' anticipation of the *winner's curse:* the winning bidder learns that other bidders had lower types, and hence that w_i is likely an overestimate of $\bar{w}$. The second term in (1) represents i's information rent, which captures the extent to which bidder i "shades" his true type in equilibrium. Information rents are greater when bidders' types are more dispersed, while the severity of the winner's curse depends on the distribution of w_i.

In practice, DOE does not sell crude using a fixed price auction format. Instead, DOE uses a (pay-as-bid) *basis auction.* In a basis auction, bids express a willingness to pay in terms of a price adjustment factor (PAF) relative to the index price W. If i wins the auction after bidding d per barrel, he is obligated pay $W + d$ per barrel to DOE upon delivery. Proposition 2 characterizes a pure-strategy Bayes Nash equilibrium of this game.

Proposition 2. *Let $y_i = \max_{j \neq i} u_j$.The N-tuple of strategies $(\delta_1(\cdot), \ldots, \delta_N(\cdot))$, where*

$$\delta_i(x) = E[y_i | y_i \leq x], \tag{2}$$

is an equilibrium of the basis auction.

Equation (2) is familiar as the equilibrium bidding strategy in the symmetric equilibrium of a first price auction with independent private values. Importantly, (2) does not depend on bidders' private information w_i or on the realization of the mean signal $\bar{w}$. For a fixed number of bidders, this difference translates into greater revenue for the seller, and to greater surplus overall.

Proposition 3. *DOE's expected revenue is greater in the basis auction than in the fixed price auction. Moreover, the expected total surplus is larger.*

The basis auction improves seller revenue for two reasons. First, because the winner of the basis auction is the bidder with the highest cost basis u_i rather than the highest type s_i, the basis auction creates greater total surplus in expectation.[8] Second, because bidders in the basis auction have a smaller amount of relevant information, they earn lower rents.

Proposition 3 offers one rationale for why DOE may prefer a basis auction to a fixed price auction. However, it is not clear that the main purpose of the SPR release is to raise revenue for DOE (Congressional Research Service, 2020), or to maximize surplus total surplus. A second rationale for why DOE may prefer a basis auction relates to efficiency: under the maintained assumption of bidder symmetry, the bidder with the highest cost basis u_i is bound to win the basis auction.[9] DOE may value this form of efficiency for at least two reasons. Because oil refiners face capacity constraints, the stabilizing effect of SPR sales on oil prices may be stronger in the case that SPR crude is allocated to refiners with greater opportunity costs. Second, divorcing allocations from the private signals w_i eliminates the incentive for costly information gathering, which could otherwise induce a wasteful "arms race."

Price adjustments for quality In practice, DOE further adjusts prices for crude oil quality. Specifically, DOE adjusts final payments in accord with the degree to which the API gravity and sulphur content of crude taken at delivery differ from levels published in the Notice of Sale. I elaborate on the significance of crude oil quality below.

Endogenous default In the fixed price auction, winning bidders are incentivized to default if W is lower than $\beta_i(w_i/n + u_i) - u_i$ at delivery. This incentive could be strong as SPR sales are typically held in the midst of rapid price increases when market and/or geopolitical volatility are elevated, which may be times when forecasts are especially noisy. It is clear that there is no such incentive in the basis auction. Because $\delta(u_i) \leq u_i$, the winning bidder always prefers delivery of SPR crude to his outside option on the delivery date.

[8] Efficiency need not hold in the case of bidder asymmetry (Maskin and Riley, 2000). In SPR sales, the distinction between oil refiners and marketers (discussed below) is potentially an important source of bidder asymmetry.

[9] In the fixed price auction this is not necessarily the case, since a bidder's type also depends on his signal w_i of the common value.

Accounting for the possibility of default could affect the comparison between the fixed price and basis auctions.[10] If SPR contracts are viewed as options rather than obligations, bidding in the fixed price auction could become more aggressive, increasing the seller's revenue conditional on sale but also increasing the incentive to default ex post, potentially reducing revenue. In practice, it may be possible to deter default via punishment. Under current DOE guidelines, defaulting bidders incur severe penalties.[11] Because bidders are long-lived and deep-pocketed in comparison to the value of a one-time crude oil purchase, such penalties may be sufficient to deter default.[12] Moreover, existing penalties are designed for a basis auction regime, in which the model implies incentives to default are absent. In a fixed price regime, DOE could impose different and stronger penalties.

Hedging with forward contracts In reality winning bidders might choose to hedge their purchases using forward contracts. If all bidders in the fixed price auction were required to hedge, and if hedging contracts were perfect and costless, the fixed price auction would be equivalent to a basis auction. Conversely, hedging renders the basis auction more similar to a fixed price auction. Thus, the effects of the basis auction emphasized in this section are most salient when forward contracts are costly to obtain or otherwise imperfect.

Other extensions In general multiunit auction environments in which bidders have heterogenous demands, declining marginal valuations, and are possibly endowed with private information, it is typically impossible to characterize bidding equilibria outside of a small number of parametric models (Rostek and Yoon, 2023). Even under the assumption of independent private values (IPV), relatively little is known about equilibrium properties of the pay-as-bid auction in particular despite its wide usage in sovereign debt auctions and other settings (Hortacsu and McAdams, 2018; Pycia and Woodward, 2025).

2.1 Empirical implications

The conclusions of the previous section depend crucially on the assumption that the price index used to determine bidders' final payments coincides with the common component of

[10] Early papers analyzing the effect of ex post withdrawal or default in auctions include von Ungergn-Sternberg (1991), Waehrer (1995), and Zheng (2001). Empirical contributions include Asker (2000) and work on delays and cost overruns in highway procurement (Lewis and Bajari, 2011; Decarolis, 2014).

[11] 1% of the contract price per day until a secondary sale can occur and 100% of any losses incurred by DOE. Defaulting bidders are further threatened with debarment from future sales.

[12] In current practice, default can occur without penalty in some pre-specified circumstances (for example, due to unexpected default of a subcontractor) or at the discretion of the contracting officer. The severity of a penalty must be judged in conjunction with the likelihood that an exemption would be granted.

valuations. If this assumption is violated, the basis auction may retain a common values element.

The price indices used to determine final payments in SPR sales may be imperfect for a number of reasons. First, transportation cost shocks particular to accessing SPR facilities could differentially raise the price of SPR crude relative to other sources of crude for all potential buyers. Second, price adjustments for API gravity and sulphur content may not fully address payoff-relevant variation in crude oil quality, which is not fully known until delivery.

To incorporate this possibility into the model, I now assume that bidders' valuations take the form $p_i = W + Z + u_i$, where Z represents a (possibly non-contractible) differential between the index price and the common component of bidders' valuations. In addition to u_i and w_i, each bidder privately observes an unbiased estimate z_i of the true but unknown ex post differential Z. The signal z_i is drawn independently across bidders from a distribution F_z having log-concave density f_z with support on a bounded interval $[\underline{z}, \bar{z}]$. z_i is independent from w_i and u_i. If final payments are indexed to W rather than $W + Z$, the basis auction format will not fully eliminate the common values component from the auction.

Proposition 4. *Define $\tilde{s}_i = z_i/N + u_i$. Let $y_i = \max_{j \neq i} \tilde{s}_j$. The N-tuple of strategies $(\delta_1(\cdot), \ldots, \delta_N(\cdot))$, where*

$$\delta_i(x) = E[\bar{z} + u_i | \tilde{s}_i = x, y_i \leq x] - E[x - y_i | \tilde{s}_i = x, y_i \leq x], \tag{3}$$

is an equilibrium of the basis auction.

Proposition 4 provides the primary testable implication of the model: the basis auction is a private values auction if and only if Z is identically zero. The (potential) benefits of the basis auction suggested above may be small in magnitude if the price index W is not significantly different from the common component of bidders valuations $W + Z$. Contrasting (2) with (3), it is clear that z_i constitutes an additional source of information rents and that bidders must now account for the winner's curse. Moreover, the bidder with the highest private cost basis u_i is no longer certain to win the auction.

The next result establishes that bids are decreasing in the variance of Z.

Proposition 5. *Let $\tilde{F}_z$ be a mean-preserving spread of F_z. Bids in the basis auction are lower under $\tilde{F}_z$ than under F_z.*

This result provides a second testable implication of the model: bid levels should be negatively correlated with instruments that shift the second moment of the common component of bidders' valuations $W + Z$ separately from W. In other words, variables that make payoffs less certain for all bidders apart from their effect on the oil price index used for price adjustment should correlate with lower bid levels in the basis auction.

Section 4 presents two alternative testing strategies motivated by these observations. Evidence in favor of private values suggests that the price indices used by DOE eliminate all empirically relevant sources of common values. Evidence of private information instead suggests that DOE could benefit from an alternative auction design, such as one with different price adjustment criteria. Before describing the testing strategies in detail, I provide additional background on SPR auctions and introduce the bidding data.

3 The 2022 Emergency Drawdown

In 2022, United States initiated seven large sales (or "drawdowns") from the SPR. In the remainder of the paper, I use bidding data from these sales to test the main prediction of the stylized model: if the price indices used to determine final payments are sufficiently similar to the common component of bidders' valuations, then firms' observed bidding strategies should be indistinguishable from a private values model. Before presenting the main results, I provide further background on SPR sales and introduce the data.

3.1 Background

The SPR was created by the 1975 Energy Policy and Conservation Act (ECPA) in response to the 1973 oil crisis. Currently, the SPR stores up to 713.5 million barrels (MMbbl) of crude oil at four distinct locations along the US Gulf Coast, shown in Figure 1. Each location stores a combination of sweet (i.e., low sulfur) and sour (i.e., high sulfur) crude oils in underground salt caverns. Locations differ in their capacities, crude oil characteristics, access to pipelines and shipping terminals, and proximity to major Gulf Coast refineries.

The main purpose of the SPR is to "protect the U.S. economy from severe petroleum supply interruptions" by distributing stockpiled oil in times of crisis (DOE, 2024). Prior to 2015, SPR sales occurred rarely and were modest in size.[13] Beginning in 2015, Congress mandated a series of sales in order to raise revenue for budgetary purposes. These sales were dwarfed in size by a series of seven emergency sales totaling 180 MMbbl conducted in 2022 in response to oil market disruptions associated with the Russia-Ukraine conflict. I focus on these seven sales; detailed bidding data for past sales is publicly available from DOE.

[13] Notable "emergency" releases occurred after Operation Desert Storm (17.3 MMbbl), Hurricane Katrina (11 MMbbl), and the 2011 Libyan crisis (30.6 MMbbl).

Figure 1: SPR System Map

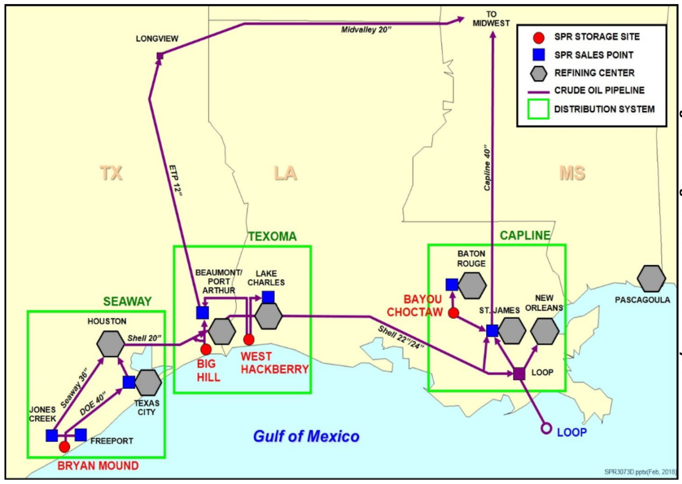

Source: DOE.

3.2 The mechanics of SPR sales

Rules governing sales from the SPR were first promulgated by DOE in 1983.[14] Each sale begins with the issuance of a Notice of Sale. In the Notice of Sale, DOE stipulates target volumes of sweet and sour crude oil to be sold from designated storage sites. Figure 2 provides an example sales specification taken from the third sale that occurred in 2022. In this sale, which occurred on June 1, the DOE sought to release one "stream" of sweet crude and one stream of sour crude from the Bayou Choctaw storage site near Baton Rouge, LA. For the sweet crude "master line item" (MLI), 1,100 thousand of barrels (Mbbl) were offered for delivery between June 21 and June 30. For the sour crude MLI, 6,000 Mbbl of sour crude were offered for delivery between July 1 and August 15. The rate at which crude oil can be transferred from storage to particular "points of sale" is limited by physical transfer capacity and scheduling constraints. The "delivery line items" (DLIs) in Figure 2 indicate that for the sweet crude MLI, there were no constraints on pipeline delivery at the St. James point of sale, but pipeline deliveries to the Baton Rouge (Bourre) point of sale were limited to 225 Mbbl of the total 1,000 Mbbl MLI quantity. For the sour crude MLI, the same two pipelines DLIs were available for most of the 6,000 Mbbl MLI quantity, but 800 Mbbl was only made available for vessel delivery (e.g., to oil barges).

A bid specifies a quantity demanded and price for delivery of oil at a particular point of sale. Firms are permitted to submit multiple bids. However, because of the logistical costs inherent to the physical transfer of oil, the minimum allowed bid size is large relative to the total quantity available (typically 225-350 Mbbl). For each MLI (crude oil stream),

[14] https://www.spr.doe.gov/reports/SSPs/PART625.html

Figure 2: Notice of Sale Example

Crude Oil Stream MLI	DLI - Mode of Delivery	Delivery Period	MLI Qty (MB)	MBD Avg	DLI Qty (MB)	MIN Qty (MB)
Bayou Choctaw Sweet 007		6/21-6/30	1,100			
	DLI-A Pipeline St. James Pipeline	6/21-6/30		300	1,100	350
	DLI-H Pipeline Baton Rouge (Bourre)	6/21-6/30		200	225	225

Crude Oil Stream MLI	DLI - Mode of Delivery	Delivery Period	MLI Qty (MB)	MBD Avg	DLI Qty (MB)	MIN Qty (MB)
Bayou Choctaw Sour 008		7/1-8/15	6,000			
	DLI-A Pipeline St. James Pipeline	7/1-8/15		300	4,300	350
	DLI-B Vessel Saint James Dock	7/1-8/15		200	800	250
	DLI-H Pipeline Baton Rouge (Bourre)	7/1-8/15		200	900	225

Source: DOE Notice of Sale (Emergency Drawdown No. 2B, May 24, 2022).

the auction is cleared as a pay-as-bid (discriminatory) auction subject to DLI (point of sale) constraints. Bids are fulfilled from high price to low price until the total MLI quantity is exhausted, unless the supply of a particular DLI is exhausted (in which case any remaining MLI quantity is allocated other DLIs). In the previous example, if the high bidder for Bayou Choctaw Sweet requested 225 Mbbl at Baton Rouge (Bourre), leaving no remaining supply, then any further bids for delivery at Baton Rouge (Bourre) would be discarded and the remaining 875 Mbbl could only be sold to bidders requesting delivery at St. James. On the other hand, if the high bidder requested 1,100 Mbbl at St. James, no further supply would be available for bidders requesting delivery at Baton Rouge (Bourre).

As discussed above, an important feature of the SPR auction is that bids are not expressed in fixed prices, but rather as a basis relative to an index price (referred to as a price adjustment factor or PAF). For sweet crude, the index price is the average over the five days surrounding delivery of the WTI Houston Weighted Average Month 1, Houston Close (or "WTI Houston") oil price index; for sour crude, the Mars Weighted Average Month 1, Houston Close (or "Mars") oil price index is used. The window for taking delivery of purchased crude typically does not begin until a few weeks after the auction, and can last more than a month (as in the Bayou Choctaw example above). During this time, the index price can move considerably. Table 1 shows the average value of WTI Houston and Mars crude oil price indices during the delivery window in comparison to the "base reference price" (BRP) listed in the Notice of Sale.[15] After the first sale, the sweet and source crude price indices

[15]The reported crude oil indices correspond to the Bloomberg tickers USCRMEHC (Crude Oil WTI Houston) and USCRMARS (Crude Oil Mars). These indices are not identical to the specific price indices

were 12% higher than the BRP on the average day during the delivery window. More often, prices fell. After the fourth sale, price indices were more than 25% lower on average.

One notable feature of the auction environment absent from the model in Section 2 is a reserve price fixed at 95% of DOE's estimate of the market value of crude. DOE reserves the right to reject bids that are "below 95 percent of the sales price, as estimated by the Government, of comparable crude oil being sold in the same area at the same time." I describe this reserve price as *nominal* because bids below it may still be accepted "if the Contracting Officer determines such action is necessary to achieve SPR crude oil supply objectives and such offered prices are reasonable." As explained below, the nominal reserve price appears not to have been binding during the 2022 SPR drawdown.

3.3 Description of bidding data

This section briefly describes patterns observed in the bidding data. To begin, Table 2 presents summary statistics on participation, demand, and winning bids for each DLI and MLI. On average, 10.9 bidders submitted 3.7 bids each for a given sweet crude MLI, often submitting bids on multiple DLIs at once. The highest bid was $1.19 greater than the index price and $4.08 greater than the lowest winning bid. The lowest winning bid was $2.89 less than the index price, but $2.24 greater than the nominal reserve. Clearing prices were generally higher for pipeline delivery in comparison to vessel delivery. Similar patterns are present for the sour crude sales. Two key differences are that auctions for sour crude had fewer bidders on average, and higher prices relative to the corresponding price index.

Domestic oil refiners were by far the largest purchasers of crude oil in the 2022 SPR sales. 74% of the crude oil sold in the 2022 SPR releases was purchased by one of the four largest US refiners (Marathon, Valero, ExxonMobil, and Phillips 66) or by Motiva (the Saudi Aramco subsidiary which owns the massive Port Arthur Refinery in Texas). Most of the remaining crude oil was purchased by widely known commodities trading firms.

Interestingly, more than 15.2% of bids by quantity were submitted below the nominal reserve price, indicating that bidders may not have expected the reserve to be enforced. Consistent with this belief, there were no auctions in which the DOE rejected otherwise competitive bids for being lower than the reserve price. However, the vast majority of these low bids were unsuccessful. Only 1.4% of accepted bids were below the nominal reserve.

used by DOE in practice, which are published exclusively by Argus Americas Crude and could not be obtained for this study. It is reasonable to expect that these series are highly correlated.

4 Tests of common and private values

In this section I aim to test whether bidders in the SPR auction bid as if they compete in a private values environment, as suggested by the stylized model in Section 2.

Exclusive tranches Empirical methods for auctions typically assume that the unit of observation is a single, independent auction with a fxied number of potential and active bidders. The latter assumption can be violated in SPR sales due to DLI constraints: with respect to a single DLI, competition may change discontinuously with quantity.[16]

To illustrate, we can again consider the case of Bayou Choctaw Sweet in Figure 2. Bidders who submitted bids for the St. James DLI faced more competition for quantities above than below 875 Mbbl because quantities in excess of 875 Mbbl were also contested by bidders who submitted bids for the Baton Rouge (Bourre) DLI.

To abstract from complications arising from DLI constraints, the primary unit of observation for the empirical analysis is an *exclusive tranche*. An exclusive tranche is defined as a portion of crude oil that accessible only to bidders on a particular DLI. In the example, there is one exclusive tranche of 875 Mbbl associated with St. James, but no exclusive tranche associated with Baton Rouge (Bourre). For Bayou Choctaw Sour, also shown in Figure 2, there were three exclusive tranches: one for each DLI. In total, there were 54 exclusive tranches (an average of 2.16 per MLI per sale). Each exclusive tranche can be viewed as a distinct pay-as-bid auction with a fixed number of bidders. Restricting attention to exclusive tranches involves little loss of data: 89% of crude oil made available during the 2022 SPR drawdown and 93% of revenue can be linked to an exclusive tranche.[17]

4.1 Test of common values

Few formal tests of private information have been developed for standard multi-unit auctions, much less a nonstandard multi-unit setting with high minimum bid quantities and within-auction capacity constraints.[18] In this section I consider a reduced form test of common

[16] A similar effect occurs in competitive wholesale electricity markets due to transmission congestion.

[17] Revenue calculation based on base reference prices. In principle one might construct additional "residual" tranches that span multiple DLIs. For example, the 225 Mbbl contestable by bidders on both St. James and Baton Rouge (Bourre) in the Bayou Choctaw Sweet example might be viewed as a tranche with a number of bidders that depends on residual demand from Baton Rouge (Bourre) bidders after the first 875 Mbbl is exhausted. I do not pursue this approach because the number of bidders is endogenous in this case.

[18] Hortaçsu and Kastl (2012) exploit the availability of an instrument shifting some bidders' information in the context of Canadian treasury auctions. Bonaldi et al. (2024) estimate a parametric model of a uniform price auction in which bidders have both private and common values, enabling a decomposition exercise. I lack the data to implement the test of Hortaçsu and Kastl (2012), while the lack of convenient and flexible parametric models for pay-as-bid auctions precludes an approach similar to Bonaldi et al. (2024).

values developed by Nyborg et al. (2002) and Bjønnes (2001).

Proposition 5 implies that bid levels in the basis auction are decreasing in the uncertainty of the residual common value Z whenever indexing fails to induce a private values auction. Nyborg et al. (2002) argue more generally that bidders in pay-as-bid auctions with common values tend to reduce the level and steepen the slope of their bids as uncertainty becomes more severe.[19] These qualitative predictions are founded on Ausubel (2004)'s notion of the "champion's plague," which generalizes the winner's curse to multiunit auctions.[20] Evidence of correlation between observed bids and uncertainty is therefore suggestive of common values.

A limitation of the Nyborg et al. (2002) framework is that similar qualitative predictions can be generated by bidder risk aversion. Bjønnes (2001) suggests that exogenous variation in participation (absent form the empirical setting of Nyborg et al. (2002)) can be used to distinguish between private information and risk aversion in this context, because increased participation exacerbates the champion's plague but dilutes (ex post) risk.

In the case of SPR auctions, one convenient measure of uncertainty is the average time between the auction date and delivery. A second is the CBOE Crude Oil Volatility Index (or OVX).[21] The latter measure closely proxies for the uncertainty of oil prices but remains informative if the specific price indices used by DOE are imperfect. The former also correlates with oil price uncertainty but more plausibly accounts for sources of uncertainty less directly related to oil prices (for example, if transportation costs are less precisely known over longer time horizons). Bid levels can be measured using a firm's quantity-weighted average PAF. The steepness of bid slopes is captured by intra-bidder bid dispersion, measured by the quantity-weighted standard deviation in PAFs across a firm's bids.

Table 4 summarizes regressions of bid levels and dispersion on both measures of uncertainty, after controlling for other key observables. As in Nyborg et al. (2002) and Bjønnes (2001), observations are weighted by the measure of uncertainty, and two specifications for the bid level are included: one using bidder-tranche level average bids, and one using the tranche-level average of bidders' within-tranche average bids.[22] Following Bjønnes (2001), I control for bidder participation. However, because bidder participation may be endogenous, I instrument for it using the location of the storage site (see Figure 1).[23] The identifying

[19] An additional prediction of Nyborg et al. (2002) is that bidders will reduce their maximum quantity demanded if the seller imposes a reserve price. I do not evaluate this prediction.

[20] The winner's curse refers to the idea that the winner of a single unit auction learns that other bidders had lower signals ("bad news"); the champion's plague refers to the idea that winning relatively more units reveals relatively worse news.

[21] See, e.g., Chen et al., 2018 for a discussion of this index.

[22] The latter helps to address bias that may arise from correlation in the error terms within an auction.

[23] Bjønnes (2001) does not instrument for bidder participation.

assumption is that location affects the intensity of competition, but not bidders' valuations. First stage regression results are reported in Table 3; the F-statistic is 8.11.

If a common values component is present after indexing bids to oil prices, we would expect to observe a negative relationship between bid levels and uncertainty. In Table 4, the coefficients on uncertainty in columns (1), (2), (4), and (5) are all positive. Thus, bid levels appear to increase when uncertainty is greater. More specifically, the point estimates imply that shifting the delivery window back by a week increases average bids by about \$0.20 per barrel. An increase of similar magnitude is implied by a 5-10 point increase in OVX. These results can be viewed as inconsistent with the presence of residual private information, although the effects are small in magnitude and none is statistically significant. The lone significant effect from uncertainty (at the 10% level) indicates that bid dispersion increases with oil price volatility, which could be consistent with private information. However, the corresponding effect of time until delivery on bid dispersion is negative. Taken together, these results do not provide strong evidence in favor of common values.

Bjønnes (2001) finds a negative association between participation and bid levels, and interprets this as evidence that private information is quantitatively more important than risk aversion. Here I find a positive effect: an additional bidder increases average bid levels by \$0.68 per barrel. Likewise, I find a negative association between participation and bid dispersion. These results, although imprecise, are suggestive of the opposite conclusion: that private information (if it is present) is not of first-order importance.

4.2 Test of private values

In this section I consider a direct test of independent private values (IPV) due to Hickman et al. (2021). The central idea underlying this test is that, under the null hypothesis of IPV with no unobserved heterogeneity, there should be no residual correlation in bids after controlling for the number of bidders N and any other observables.[24]

In a single-unit context, this idea can be implemented by regressing observed bids on the mean of within-auction rival bids, observables, and a flexible function of N. IPV is rejected if the coefficient on within-auction rival bids is different from zero. This idea can be extended to the context of a multiunit auction by restricting attention to quantity levels that are common to all submitted bids. For example, if all bidders are willing to accept a single unit of the good, one might use the highest price bid submitted by each bidder.

Simulation results and empirical applications in Hickman et al. (2021) suggest that this test has favorable power for single-unit auctions even in small samples. It rejects not only

[24] The test does not necessarily require an assumption of exogenous participation.

common values, but also various alternative private values models such as the affiliated private value (APV) model, the conditional independent private values (CIPV) model, and IPV models with auction-level unobserved heterogeneity. In Appendix B, I verify the finite-sample power of the test against CIPV and pure common values using simulated bidding data from a uniform price auction model calibrated to SPR sales.

A further practical complication that arises in the context of SPR sales is that bidders are allowed to stipulate a minimum bid quantity that may be much larger than the maximum quantities demanded by rival bidders. To address this issue, I implement the test separately for all exclusive tranches and for the subsample of exclusive tranches in which no such bids were submitted. In the former case, I retain bids with high minimum quantities for the purpose of calculating N but exclude them from the estimation sample and from the calculation of the mean of rivals' bids.

Table 5 reports the test results. As in Richert (2024b), each specification includes a cubic polynomial of N. Focusing on the first column, the coefficient on the mean rival bid is almost precisely zero. Thus, the test fails to reject the null of IPV in the subsample of auctions without large minimum bid quantities. A similar result is shown in the second column, which corresponds to the full sample of auctions. Together with the results above, these findings provide suggestive evidence that the SPR auction mechanism induces bidding strategies consistent with independent private values.

5 Gains from price adjustment

In this section I quantify a stylized equilibrium model of a multiunit auction intended to resemble the 2022 SPR drawdown. I use the model to benchmark the impacts of price adjustment for different information structures. While only illustrative, the results show how price adjustment could have significantly increased DOE's expected revenue.

Vives (2011) studies a parametric model of a uniform price auction with an information structure that features both private and common values. I build on this model. First, I calibrate the model parameters using observed bidding data as well as contemporaneous oil price forecasts.[25] The calibration suggests bidders' valuations would have been highly correlated in a fixed price auction because future oil prices are significantly more variable than bidders' private cost bases. Given the possibility of private information, this suggests that adverse selection (the winner's curse, or in this case the champion's plague) might have been severe, depressing revenue and leading to inefficient allocations.

[25]Because the bidding data is generated from a pay-as-bid auction, among other major simplifications, this procedure cannot be viewed as estimation.

5.1 Benchmark model

There are N bidders and Q barrels of oil for sale. Bidders' marginal valuations take the form:

$$v_i(q) = W + u_i - \lambda q$$

where $\lambda > 0$ captures inventory costs (or any other source of diminishing marginal valuations). W is the realized oil price index at delivery, which is assumed to be distributed $N(\mu_w, \sigma_w^2)$. u_i is the bidder i's private cost basis and is distributed $N(\mu_u, \sigma_u^2)$.

Bidders compete by submitting demand schedules in a uniform price auction. Prior to the auction each bidder observes a signal $s_i = W + u_i + \epsilon_i$ where $\epsilon_i \sim N(0, \sigma_\epsilon^2)$ captures bidder i's uncertainty about his own valuation. u_i and ϵ_i are mutually independent, independent of W, and independent across bidders. Bidders valuations are potentially correlated through W. For all bidders $i \neq j$, $\text{cor}(W + u_i, W + u_j) = \rho$. This structure differs from the one considered in Section 2 but encompasses the IPV model and (asymptotically) the pure common values model, along with many intermediate information structures.[26]

I focus on linear, symmetric Bayes Nash Equilibria (BNE). Vives (2011) shows that a BNE exists and is unique provided that the number of bidders N is sufficiently large relative to the dispersion of information. Equilibrium demand schedules have a simple closed form.

I maintain the assumption throughout this section that price adjustment fully eliminates the common component of valuations in DOE's pay-as-bid basis auction. Under this assumption, DOE's basis auction can be represented in the framework of Vives (2011) by a uniform price auction in which $\mu_w = \sigma_w = \sigma_\epsilon = 0$: bidders know their own valuations perfectly but only know the distribution of rivals' valuations; there is no common value.

The fixed price auction can be represented by a uniform price auction in which $\sigma_w, \sigma_\epsilon > 0$ and $\rho = \sigma_w^2/(\sigma_w^2 + \sigma_u^2)$. When $\sigma_\epsilon > 0$ bidders are uncertain about their own valuations, as in the classical common values model. Given $\sigma_w > 0$ and $\rho = \sigma_w^2/(\sigma_w^2 + \sigma_u^2)$, the common value component of valuations coincides with W.

The set of parameters necessary to compare auction formats is therefore $\theta^\tau = (\lambda^\tau, \mu_u^\tau, \sigma_u^\tau, \mu_w^\tau, \sigma_w^\tau, \sigma_\epsilon^\tau)$, where τ denotes the tranche. Except where necessary I suppress the τ superscripts.

5.2 Calibration

This section briefly describes the calibration strategy; further details are provided in Appendix C. Note that nine exclusive tranches having $N \leq 2$ observed bidders were dropped

[26] See Figure 3 along with the discussion in Section B.

from the sample used for calibration due to equilibrium nonexistence in the uniform price basis auction. I defer discussion of σ_ϵ to the next section.

Oil price distribution μ_w^τ is calibrated based on the daily average of the front-month WTI Futures Contract prices during the delivery period as evaluated at the last date prior to the sale. I adjust these prices (upwards for sweet, downwards for sour) by a constant basis reflecting differences between WTI and the oil price indices for sweet and sour crude specified in the Notices of Sale. σ_w^τ is calibrated from the above-mentioned OVX index.

Cost basis distribution I select μ_u^τ such that expected revenue in the uniform price basis auction coincides with DOE's realized revenue. I fix a single value $\sigma_u^\tau = \sigma_u$ for all tranches τ based on the mean standard deviation of purchase shares within each tranche.

Inventory costs I calibrate $\lambda^\tau = \lambda$ based on the average slope of aggregate demand in the vicinity of the clearing price after making an adjustment for competition. Intuitively, bidders submit flatter demand schedules when facing more competition; therefore, the slope of aggregate demand must be normalized across auctions with N.

5.2.1 Model validation

Table 6 reports sample moments from the bidding data alongside their expected analogues in the calibrated uniform price basis auctions. By construction, revenue and quantity dispersion are similar in the model and in the sample. Revenue matches exactly, while dispersion differs slightly due to undersubscription of some exclusive tranches in the data. One informative moment not targeted in the calibration is the bid-to-cover ratio, defined here as the total quantity demanded at the minimum observed price tick relative to the tranche quantity. The average bid-to-cover ratio in the sample was 3.04 versus 3.06 for the calibrated model.

5.3 Comparison of auction formats

Because the calibrated dispersion of oil prices σ_w is typically large relative to the calibrated dispersion in opportunity costs σ_u, ρ is close to one for every tranche τ. On average across tranches the calibrated correlation is 0.73 (95% bootstrap CI: $[0.71, 0.75]$), suggesting that fixed price auctions would more closely resemble common than private values auctions.

A key difficulty in analyzing the fixed price auction is that the bidders' valuation uncertainty σ_ϵ (more conveniently, bidders' posterior uncertainty $\sigma_{W+u|s}$) is not easily quantified.

The comparison to basis auctions depends crucially on this parameter: the expected revenue gain from the basis auction format increases in bidders' posterior uncertainty.

This can be seen as follows. If bidder i directly observes his own cost basis u_i, $\sigma_{W+u|s}$ exclusive quantifies i's residual uncertainty about the oil price index W after observing s_i. When $\sigma_{W+u|s} = 0$, there is no residual uncertainty (bidder i knows W with certainty), and the information structure simplifies to conditional independent private values (CIPV). In this case, the seller's expected revenue in the fixed price and basis price auctions is the same. For $\sigma_{W+u|s} > 0$ uncertainty about W gives rise to adverse selection, leading bidders to submit steeper demand curves and pushing equilibrium prices downwards. Hence, the basis auction improves revenue even for small amounts of uncertainty about W for reasons similar to Riley (1988). When $\sigma_{W+u|s}$ is sufficiently large, equilibrium demand curves become nearly vertical and the clearing price diverges to $-\infty$ (Vives, 2011).[27]

The key question is whether expected revenue gains from the basis auction are economically significant for plausible levels of posterior uncertainty $\sigma_{W+u|s}$. To address this question, I calculate the threshold uncertainly level $\sigma^*_{W+u|s}$ in each tranche required for the basis auction to improve expected revenue by \$1 per barrel.[28] This benchmark is notable because expected revenue gains of \$1 per barrel would imply total gains of more than \$150 million in the 2022 SPR drawdown (vs. \$17 billion in total), a significant amount. I also calculate $\sigma^\infty_{W+u|s}$, the supremum of posterior uncertainty consistent with BNE given N.

A ratio $\sigma^*_{W+u|s}/\sigma_w = 1$ implies that the revenue gain from the basis auction is equal to \$1 per barrel when bidders receive uninformative signals about W (i.e., bidder i's posterior belief is $\sigma_{W+u|s} \sim N\left(\mu_w + u_i, \sigma_w^2\right)$). A ratio $\sigma^*_{W+u|s}/\sigma_w < 1$ implies that the revenue gain from the basis auction is equal to \$1 per barrel or more whenever bidders have posterior uncertainty $\sigma_{W+u|s} \in (\sigma^*_{W+u|s}, \sigma^\infty_{W+u|s})$.[29] Across tranches in the calibrated model I find an average ratio of $\sigma^*_{W+u|s}/\sigma_w = 0.44$ (95% bootstrap CI: $[0.37, 0.50]$) and normalized interval width $(\sigma^\infty_{W+u|s} - \sigma^*_{W+u|s})/\sigma_w = 0.40$ (95% bootstrap CI: $[0.36, 0.43]$), implying significant revenue gains over a wide range of information structures. This includes information structures in which bidders have high precision signals and adverse selection is modest, as well as those in which bidders know little more than the prior and adverse selection is severe.

[27] This pathology is well-known in the literature. To avoid it, it is conventional to assume that the aggregate supply curve is elastic (or stochastic) either at the discretion of the seller or due to a fringe of price-taking bidders ("noise traders"). Neither of these assumptions is well suited to the case of SPR auctions, where quantities are fixed in advance and all bidders are sophisticated.

[28] This calculation holds tranche participation N_τ fixed. The analysis suggests that a switch from basis to fixed price auctions is more likely to induce exit than entry, further reducing expected revenue.

[29] Above $\sigma^\infty_{W+u|s}$ the model does not yield a clear prediction because a linear BNE does not exist. Other types of equilibria may exist: for example, price-taking equilibria (PTE) in which bidders learn from prices but do not account for price impact exist under less stringent conditions on N.

If W is difficult to predict even for sophisticated market participants, and if the pay-as-bid and uniform price auctions are expected to deliver similar outcomes empirically, then these results are suggestive that the use of basis auctions plausibly delivered a revenue improvement on the order of tens of millions of dollars (or more) in 2022.

6 Conclusion

This paper investigates the role of price adjustment in SPR release auctions. First, I use an illustrative model to show that price adjustment improves efficiency, reduces information rents, and reduces the risk of default. If the common component of bidders' valuations is exactly equal to the price index, bidders in the model behave as if they compete in a private values environment despite being endowed with private information about future oil prices. Empirically, I fail to find strong evidence of common values in bidding data from the 2022 SPR drawdown, consistent with the implications of the model. This suggests that price adjustment of the form used by DOE is an effective policy in the context of SPR sales. Moreover, the apparent success of price adjustment in SPR sales, which broadly resemble other forms of upstream contracting in oil markets, helps to complement existing theoretical explanations for why price adjustment is often observed in long-term contracts – an important question which has proven difficult to address empirically.

References

Asker, J. (2000) “Bidding up, buying out and cooling-off: an examination of auctions with withdrawal rights,” *Journal of Economic Theory*, 16, 585–611.

Ausubel, L.M. (2004) “An Efficient Ascending-Bid Auction for Multiple Objects,” *American Economic Review*, 94 (5), 1452–1475.

Ausubel, L.M., P. Cramton, M. Pycia, M. Rostek, and M. Weretka (2014) “Demand Reduction and Inefficiency in Multi-Unit Auctions,” *Review of Economic Studies*, 81, 1366–1400.

Bhattacharya, V., A. Ordin, and J.W. Roberts (2022) “Bidding and Drilling under Uncertainty: An Empirical Analysis of Contingent Payment Auctions,” *Journal of Political Economy*, 130 (5), 1319–1363.

Bjønnes, G.H. (2001) “Winner’s Curse in Discriminatory Price Auctions: Evidence from the Norwegian Treasury Bill Auctions,” *SIFR Research Report Series*, 3, Stockholm Institute for Financial Research.

Bonaldi, P., S. Téllez, and C. Gómez (2024) “Quantifying Private Information about Fundamentals in Treasury Auctions.”

Bulow, J. and P. Klemperer (2002) “Prices and the Winner’s Curse,” *The RAND Journal of Economics*, 33 (1), 1–21.

Che, Y.-K. and J. Kim (2010) “Bidding with Securities: Comment,” *American Economic Review*, 100, 1929–1935.

Chen, H., L. Liu, and X. Li (2018) “The predictive content of CBOE crude oil volatility index,” *Physica A*, 492, 837–850.

Congressional Research Service (2020) “The Strategic Petroleum Reserve: Background,Authorities, and Considerations,” Greenley, H.L., R46355.

Decarolis, F. (2014) “Awarding Price, Contract Performance, and Bids Screening: Evidence from Procurement Auctions,” *American Economic Journal: Applied Economics*, 6 (1), 108–132.

DeMarzo, P.M, I. Kremer, and A. Skrzypacz (2005) “Bidding with Securities: Auctions and Security Design,” *American Economic Review*, 95 (4), 936–959.

DOE (2024) “Strategic Petroleum Reserve Factsheet,” US Department of Energy. URL: https://www.energy.gov/sites/default/files/2024-09/strategic

Elsinger, H. and C. Zulehner (2007) “Bidding Behavior in Austrian Treasury Bond Auctions,” *Monetary Policy & the Economy, Oesterreichische Nationalbank (Austrian Central Bank)*, 109–125, Issue 2.

Goeree, J.K. and T. Offerman (2003) "Competitive bidding in auctions with private and common values," *The Economic Journal*, 113, 598–613.

Goldberg, V.P. (1985) "Price Adjustment in Long-Term Contracts," *19855 Wis. L. Rev. 527*.

Goldberg, V.P. and J.R. Erickson (1987) "Quantity and Price Adjustment in Long-Term Contracts: A Case Study of PetroleumCoke," *The Journal of Law & Economics*, 30 (2), 369–398.

Hafner, S. (2023) "Risk Aversion in Share Auctions: Estimating Import Rents from TRQs in Switzerland," *Quantitative Economics*, 14 (2), 419–470.

Hansen, R.G. (1985) "Auctions with Contingent Payments," *American Economic Review*, 75 (4), 862–865.

Hickman, B., T. Hubbard, and E. Richert (2021) "Testing independent private values versus alternative models in auction models."

Hortaçsu, A. (2002) "Bidding Behavior in Divisible Good Auctions: Theory and Evidence from the Turkish Treasury Auction Market," *Working Paper*.

Hortaçsu, A. and J. Kastl (2012) "Valuing Dealers' Informational Advantage: A Study of Canadian Treasury Auctions," *Econometrica*, 80 (6), 2511–2542.

Hortacsu, 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.

——— (2018) "Empirical Work on Auctions of Multiple Objects," *Journal of Economic Literature*, 56 (1), 157–184.

Joskow, P.L. (1988) "Price Adjustment in Long-Term Contracts: The Case of Coal," *The Journal of Law and Economics*, 31 (1), 47–83.

Killian, L. and X. Zhou (2020) "Does Drawing Down the U.S. Strategic Petroleum Reserve Help Stabilize Oil Prices?" *Journal of Applied Econometrics*, 35 (6), 673–691.

Kong, Y., I. Perrigne, and Q. Vuong (2022) "Multidimensional Auctions of Contracts: An Empirical Analysis," *American Economic Review*, 112 (5), 1703–1736.

Lewis, G. and P. Bajari (2011) "Procurement Contracting With Time Incentives: Theory and Evidence," *The Quarterly Journal of Economics*, 126 (3), 1173–1211.

Mariño, E.A.G. and D. Marszalec (2023) "Strategic supply management and mechanism choice in government debt auctions: An empirical analysis from the Philippines," *Journal of Banking & Finance*, 154, 106945.

Maskin, E. and J. Riley (2000) "Asymmetric Auctions," *Review of Economic Studies*, 67 (3), 413–438.

Nachman, D. and T. Noe (1994) "Optimal design of securities under asymmetric information," *Review of Financial Studies*, 7 (1), 1–44.

Newell, R.G. and B.C. Prest (2017) "Informing SPR policy through oil futures and inventory dynamics," NBER Working Paper 23974.

Nyborg, K.G., K. Rydqvist, and S.M. Sundaresan (2002) "Bidder Behavior in Multiunit Auctions: Evidence from Swedish Treasury Auctions," *Journal of Political Economy*, 110 (2), 394–424.

Oren, S.S. and S.H. Wang (1986) "Optimal Strategic Petroleum Reserve Policies: A Steady State Analysis," *Management Science*, 32 (1), 14–29.

Pycia, M. and K. Woodward (2025) "A Case for Pay-as-Bid Auctions," Working Paper.

Razek, N., V. Galvani, R. Surya, and B. McQuinn (2023) "Can U.S. strategic petroleum reserves calm a tight market exacerbated by the Russia-Ukraine conflict?" *Resources Policy*, 86, Part B, 104062.

Richert, E. (2024a) "Indirect Inference Techniques as an Equilibrium Computation Method," Working Paper.

——— (2024b) "Quantity Commitments in Multiunit Auctions: Evidence fro Credit Event Auctions."

Riley, J.G. (1988) "Ex Post Information in Auctions," *Review of Economics Studies*, 55 (3), 409–429.

Rostek, M.J. and J.H. Yoon (2023) "Imperfect Competition in Financial Markets: Recent Developments," Working Paper.

Shaked, M. and J.G. Shanthikumar (2007) "Stochastic Orders," Springer: New York.

Skrzypacz, A. (2013) "Auctions with contingent payments–An overview," *International Journal of Industrial Organization*, 31, 666–675.

Stevens, R.B. and J.Y. Zhang (2021) "Buyer Beware: The Asymmetric Impact of the Strategic Petroleum Reserve on Crude Oil Prices," *The Energy Journal*, 42 (6), 69–90.

Teisberg, T.J. (1981) "A Dynamic Programming Model of the U. S. Strategic Petroleum Reserve," *The Bell Journal of Economics*, 12 (2), 526–546.

Vives, X. (2011) "Strategic Supply Function Competition With Private Information," *Econometrica*, 79 (6), 1919–1966.

von Ungergn-Sternberg, T. (1991) “Swiss Auctions,” *Economica*, 58 (231), 341–357.

Waehrer, K. (1995) “A Model of Auction Contracts with Liquidated Damages,” *Journal of Economic Theory*, 67, 531–555.

Zheng, C. Z. (2001) “High Bids and Broke Winners,” *Journal of Economic Theory*, 100, 129–171.

A Proofs

Propositions 1 and 4 follow directly from Proposition 1 of Goeree and Offerman (2003). The proof of Proposition 2 below closely follows the same argument.

Proposition 2

First, observe that we can write the conditional expectation in (2) as:

$$E\left[y_i|y_i \leq x\right] = \int_{\underline{\text{u}}}^{x} t \cdot \frac{f_{y_i}(t)}{F_{y_i}(x)} \cdot dt$$

Differentiating the right hand side with respect to x gives:

$$\delta_i'(x) = \left[x - E\left[y_i|y_i \leq x\right]\right] \frac{f_{y_i}(x)}{F_{y_i}(x)}$$

The term in brackets is strictly positive because the event $y_i = x$ occurs with probability zero. As $f_{y_i}(\cdot)$ is bounded above zero, it follows that $\delta_i(\cdot)$ is strictly increasing.

Suppose that bidders $2, \ldots, N$ adopt the strategy (2). Then bidder 1's expected profit is:

$$\pi_1(d) = \{u_1 - d\} F_{y_1}\left[\delta^{-1}(d)\right]$$

The derivative of $\pi_1(\cdot)$ with respect to d is:

$$\pi_1'(d) = \left\{-\left[\delta^{-1}(d) - E\left[y_1|y_1 \leq \delta^{-1}(d)\right]\right] + \{u_1 - d\}\right\} \frac{f_{y_1}\left[\delta^{-1}(d)\right]}{\delta'\left(\delta^{-1}(d)\right)}$$

Evaluating $\pi_1'(\cdot)$ at $d = \delta(x)$ yields:

$$\pi_1'(\delta(x)) = \{u_1 - x\} \frac{f_{y_1}(x)}{\delta'(x)}$$

Since we have shown that $f_{y_1}(\cdot)/\delta'(x)$ is strictly positive, this shows that $\delta(x)$ is bidder 1's unique optimal bid.

Proposition 3

We make an argument in two parts. First, we show that expected information rents in the basis auction must be smaller than expected information rents in the first price auction. This follows from the following general property of order statistics for independent random variables drawn from log-concave distributions.

Lemma 6. *Suppose that $Y_1, \ldots, Y_N$ and $X_1, \ldots, X_N$ are samples of independent random variables, with $Y_i = X_i + U_i$ for $U_i \perp X_i$ and X_i having a log-concave density. Then:*

$$E\left[Y_{(n:n)} - Y_{(n-1:n)}\right] \geq E\left[X_{(n:n)} - X_{(n-1:n)}\right]$$

where $Y_{(n:n)}$ and $Y_{(n-1:n)}$ are the nth and $(n-1)$th order statistics.

Proof. By Theorem 3.B.7 of Shaked and Shanthikumar (2007), it follows that $X \leq_{disp} X + U = Y$ (that is, Y dominates X in the dispersive order). Theorem 3.B.31 says that if $X \leq_{disp} Y$, then $V_{i:n}$ dominates $U_{i:n}$ in the stochastic order for $i \geq 2$, where $U_{i:n} = X_{i:n} - X_{i-1:n}$ and $V_{i:n} = Y_{i;n} - Y_{i-1:n}$. □

To apply the lemma, note that the expected information rent is equal to the expected difference between the first and second order statistics.

Next, we argue that the winning bidder's pseudo-valuation in the fixed price auction is lower in expectation than in the basis price auction. This is because:

$$\begin{aligned}
E\left[S_{(n:n)}\right] + \left(\frac{n-1}{n}\right) E\left[w_j | s_j \leq S_{(n:n)}\right] &= \frac{1}{n} E\left[w_i | s_i = S_{(n:n)}\right] + \left(\frac{n-1}{n}\right) E\left[w_j | s_j \leq S_{(n:n)}\right] \\
&\quad + E\left[u_i | s_i = S_{(n:n)}\right] \\
&= E\left[w_i\right] + E\left[u_i | s_i = S_{(n:n)}\right] \\
&\leq E\left[W\right] + E\left[U_{(n:n)}\right]
\end{aligned}$$

where the second equality follows from the observation that:

$$\begin{aligned}
E\left[w_i\right] &= E\left[w_i | s_i = S_{(n:n)}\right] P\left(s_i = S_{(n:n)}\right) + E\left[w_i | s_i \leq S_{(n:n)}\right] \left\{1 - P\left(s_i = S_{(n:n)}\right)\right\} \\
&= E\left[w_i | s_i = S_{(n:n)}\right] \frac{1}{n} + E\left[w_i | s_i \leq S_{(n:n)}\right] \cdot \left(\frac{n-1}{n}\right)
\end{aligned}$$

Combining the previous results gives that the basis price auction generates greater expected surplus, while reducing expected information rents. Thus, expected revenue is improved.

Proposition 5

First, expected surplus is lower under $\tilde{F}_z$ because

$$E_{\tilde{F}_z}\left[u_i | s_i = S_{(n:n)}\right] \leq E_{F_z}\left[u_i | s_i = S_{(n:n)}\right]$$

Second, information rents are larger because of Lemma 6. Thus, the seller's expected revenue is lower. In a symmetric and monotone equilibrium, it follows that bid levels are lower.

B Simulation Study

This section evaluates the power of the Hickman et al. (2021) test in small samples.

In order to generate simulated bid data, I adopt the uniform price auction model of Vives (2011). Bidders are assumed to have marginal valuations of the form $v_i(q) = \theta_i q - \lambda q$, where $\theta_i = W + u_i$. The private cost bases u_i are distributed $u_i \sim N(0, \sigma_u^2)$ and the oil price signals w_i are distributed $w_i \sim N(0, \sigma_w)$. Both u_i and w_i are assumed to be independent from each other and independent across bidders. I assume $(\lambda, \sigma_u, \sigma_w) = (100, 10, 20)$ based on typical calibrated values from Section 5.[30] Let $\tilde{\sigma}_e^2$ denote the inverse of $\frac{1}{\sigma_w^2} - \frac{1}{\sigma_w^2+\sigma_u^2}$, which implies that $Var(\theta_i|s_i) = \sigma_w$ when $\rho = \frac{\sigma_w^2}{\sigma_w^2+\sigma_u^2}$ (as in Section 5). In the exercises below I vary ρ and σ_e without restricting the posterior variance $Var(\theta_i|s_i)$.

Figure 3 depicts the information structure in the space of ρ and σ_ϵ, where ρ denotes the correlation $\mathrm{cor}(\theta_i, \theta_j) = \rho$ for all $i \neq j$. The origin corresponds to independent private values. The vertical line that begins at $(0, 0)$ encompasses IPV models in which bidders do not perfectly observe their own valuations. The x-axis, where $\sigma_\epsilon = 0$, coincides the family of conditional IPV models with increasingly greater correlation between bidder values. The limit of perfect correlation in the CIPV model is a model of pure common values in which all bidders are perfectly informed about the value of the good. The vertical line that begins at the point $(1, 0)$ encompasses classical pure common values with uncertain valuations, with the severity of the winner's curse increasing in σ_ϵ. Finally, the interior shaded region encompasses a class of mixed information structures with both private and common values.

The simulation design is as follows. To evaluate the power of the model, I consider the ability of the Hickman et al. (2021) test to reject sequences of models that correspond to rays through the origin like the dashed green line in Figure 3. For instance, the ray defined by $(t, \alpha t)$ for $\alpha = 0$ encompasses the family of CIPV models described above. Motivated by the data, I generate samples of S auctions for $S \in \{50, 100, 500\}$. The total quantity sold in each auction i is fixed at $Q = 1000$ Mbbl. The number of bidders in each simulated auction is fixed at $N = 5$, the average number observed in the data.[31]

Figure 4 presents simulated power curves for a 5% test for each sample size along the rays $\alpha \in \{0, 0.5\tilde{\sigma}_e, \tilde{\sigma}_e\}$. Even for the case of $S = 50$ auctions, which is slightly less than

[30] Quantities are denominated in thousands of barrels (Mbbl).

[31] Note that when N is small (relative to the amount of private information), a symmetric linear BNE may not exist. I drop auctions for which no such BNE exists.

the number of exclusive tranches in the data (54), the test attains rejection rates above 75% for $\rho > 0.25$ in the CIPV model. As σ_e increases for the same level of ρ, corresponding to the introduction of private information, rejection rates decrease slightly but remain high for modest deviations from IPV. One apparent limitation of the test implied by the results is that the rejection rate at $t = 0$ is around 10% for all designs, above the nominal level.[32]

C Calibration Details

Oil price distribution First I discuss the calibration of oil price uncertainty. Let o_t denote the value of the CBOE Crude Oil Volatility Index (OVX) on date t. o_t is an annualized estimate of the expected 30-day standard deviation of oil prices.[33] It is closely related to the better known CBOE VIX, except that it is constructed from options on a particular oil price ETF (the United States Oil Fund or USO) rather than options on the S&P 500 Index. For a sale on date t of crude oil of sulphur type $s \in \{\text{sweet}, \text{sour}\}$, I take

$$\sigma_w^{t,s} = p^{t,s} \cdot \sqrt{\frac{|D_t|^{-1} \sum_{s \in D_t} (s - t)}{365}} \cdot \frac{o_t}{100}$$

where D_t is the set of days in the delivery window and $p^{t,s}$ is an estimate of the daily average USO price during the delivery window. USO itself is constructed to track the price of West Texas Intermediate Light Sweet Crude. To construct $p^{t,s}$, I therefore define:

$$p^{t,0} = |D_t|^{-1} \sum_{s \in D_t} f_{t-1,s}^{WTI}$$

where $f_{t-1,s}^{WTI}$ is the price at date $t-1$ of the front month WTI futures contract corresponding to date s during the delivery window D_t, obtained from EIA. Finally, I define:

$$p^{t,s} = p^{t,0} + \delta_s$$

where δ_s is a basis between the price index corresponding to s in DOE's Notice of Sale (WTI Houston, Mars) estimated from the average difference between the front month WTI price and USCRMEHC/USCRMARS in 2022.[34]

[32] In larger sample sizes the rejection rate appears to converge to the nominal level.

[33] https://www.cboe.com/us/indices/dashboard/ovx/

[34] In particular, I obtain $\delta_{sweet} = 1.22$ and $\delta_{sour} = -3.29$.

Cost basis distribution and inventory costs In the uniform price approximation to the basis auction, the dispersion in purchase shares $X^*_{\tau i}/Q_\tau$ is

$$Var\left(X^*_{\tau i}/Q_\tau\right) = \frac{1}{Q_\tau^2}\left(1-\frac{1}{N}\right)\cdot\left(\frac{N-2}{N-1}\right)^2\cdot\left\{\lambda^{-1}\sigma_u\right\}^2$$

I therefore take $\lambda^{-1}\sigma_u$ as the sample mean of

$$\sqrt{\left(1-\frac{1}{N}\right)^{-1}\cdot\left(\frac{N-2}{N-1}\right)^{-2}\cdot Var\left(X^*_{\tau i}\right)}$$

in the sample of exclusive tranches having $N > 2$.[35]

For λ^{-1}, I observe that the slope m_τ of the quantity-normalized aggregate demand $\sum_{i\leq N} X\left(s_i, p\right)/Q_\tau$ for tranche τ in the uniform price model is given by

$$m_\tau = \frac{N}{Q_\tau}\cdot\left(\frac{N-2}{N-1}\right)\cdot\lambda^{-1}$$

I obtain a value for λ^{-1} as the point estimate of the slope of $x_{\tau k} = \frac{N}{Q_\tau}\left(\frac{N-2}{N-1}\right)b_{\tau k}$ in a weighted linear regression of $\sum_{i\leq N} X\left(s_i, p\right)/Q_\tau$ on the collection of observed bid prices $b_{\tau 1}, \ldots b_{\tau K}$ for each τ and tranche fixed effects, where the weight of price tick τk is $w_{\tau k} = \phi\left(\sum_{i\leq N} X\left(s_i, b_{\tau k}\right)/Q_\tau - 1\right)$. This procedure approximates the average slope of the aggregate demand curve in the vicinity of the clearing price, implicitly giving greater weight to tranches with more participation (which tend to have more bid ticks). The normalization of $b_{\tau k}$ by $\frac{N}{Q_\tau}\left(\frac{N-2}{N-1}\right)$ varies with the level of competition to reflect the impacts of market power and adverse selection (the winner's curse).

[35] Under IPV, no symmetric linear BNE exists for $N \leq 2$.

Table 1: Mean Delivery Index Price (DIP) vs. Base Reference Price (BRP)

		Sweet Crude			Sour Crude		
Date of Sale	Delivery Window	BRP	DIP	% Δ	BRP	DIP	% Δ
2022-03-08	24-84 days	95.15	106.99	12.44	92.60	103.76	12.06
2022-04-12	33-79 days	111.00	116.57	5.02	106.28	108.34	1.93
2022-06-01	20-75 days	113.39	110.51	-2.54	108.81	94.54	-13.11
2022-06-28	49-94 days	121.42	89.15	-26.58	114.21	85.47	-25.16
2022-08-02	45-80 days	103.68	87.39	-15.71	97.61	82.77	-15.20
2022-09-27	35-64 days	89.45	87.56	-2.12			
2022-10-25	37-67 days	89.91	77.07	-14.28	83.86	71.39	-14.87

Notes: (1) Delivery index price is calculated from daily closing prices of Bloomberg tickers USCRMEHC (Crude Oil WTI Houston) and USCRMARS (Crude Oil Mars), see main test for discussion; (2) Base reference prices are obtained from auction-specific Notices of Sale; (3) Delivery windows vary by MLI and DLI, the maximum range is recorded.

Table 2: Summary of Bidding Data (Sample Means)

	Sweet Crude			Sour Crude		
	Per DLI (Pipeline)	Per DLI (Vessel)	Per MLI	Per DLI (Pipeline)	Per DLI (Vessel)	Per MLI
Number of bidders	6.10	4.89	10.91	4.05	3.69	8.00
Number of bids	19.65	5.67	40.36	13.82	6.31	27.57
Number of winning bidders	3.30	1.11	6.00	2.36	1.23	4.36
Number of winning bids	8.55	1.25	16.45	7.71	3.20	13.86
Total quantity demanded	9906.75	2961.11	23962.27	7559.77	3880.38	16272.14
Total quantity sold	3825.00	643.75	7422.73	3927.14	1654.00	7072.14
Maximum winning PAF	0.42	-2.21	1.19	1.28	0.90	1.85
Average winning PAF	-0.76	-2.23	-0.93	0.23	0.06	-0.08
Minimum winning PAF	-2.03	-2.24	-2.89	-0.56	-0.58	-1.44
Nominal reserve PAF (5% BRP)	-5.12	-5.06	-5.13	-5.18	-5.21	-5.15

Notes: 2022 SPR drawdown bidding data. Includes all bids. (1) PAF refers to price adjustment factor; (2) Bid quantities are expressed in thousands of barrels (Mbbl).

Table 3: First Stage Regression Results

	Dependent variable:
	N
Big Hill	1.238
	(0.894)
Bryan Mound	0.453
	(0.708)
West Hackberry	1.404*
	(0.764)
Sweet Crude	1.847***
	(0.605)
Vessel DLI	0.858
	(0.541)
Tranche Size (MMbbl)	0.508***
	(0.129)
Constant	1.259*
	(0.744)
Observations	54
Adjusted R^2	0.446
F Statistic	8.110*** (df = 6; 47)

*p<0.1; **p<0.05; ***p<0.01

Notes: (1) Means and standard deviations of price adjustment factors (PAFs) are weighted by bid quantity; (2) Volatility Index (OVX) is the closing price of the CBOE Crude Oil Volatility Index on the auction date; (3) Observations are weighted by uncertainty (Days Until Delivery in (1)-(3), Volatility Index (OVX) in (4)-(6); (4) Table reports unclustered heteroskedasticity-robust standard errors.

Table 4: Correlations between Bids and Uncertainty

	Dependent variable:					
	Mean PAF		Std. PAF	Mean PAF		Std. PAF
	(1)	(2)	(3)	(4)	(5)	(6)
Days Until Delivery	0.030	0.036	−0.007			
	(0.022)	(0.025)	(0.005)			
Volatility Index (OVX)				0.023	0.042	0.014*
				(0.040)	(0.028)	(0.008)
N	0.523	0.680	−0.087	0.589	0.668	−0.061
	(0.553)	(0.504)	(0.151)	(0.514)	(0.421)	(0.127)
Sweet Crude	−2.037	−2.372**	0.139	−2.149*	−2.310***	0.084
	(1.330)	(1.127)	(0.406)	(1.237)	(0.826)	(0.359)
Vessel DLI	−2.815***	−2.720***	0.042	−2.809***	−2.701***	0.009
	(0.787)	(0.587)	(0.179)	(0.739)	(0.484)	(0.170)
Tranche Size (MMbbl)	−0.594*	−0.711**	0.109	−0.641**	−0.729***	0.091
	(0.329)	(0.303)	(0.086)	(0.308)	(0.256)	(0.076)
Constant	−3.431*	−3.997*	0.834**	−3.046	−4.020*	−0.345
	(1.938)	(2.216)	(0.362)	(2.405)	(2.056)	(0.388)
Unit of Observation	Bidder-Tranche	Tranche	Bidder-Tranche	Bidder-Tranche	Tranche	Bidder-Tranche
Observations	268	54	268	268	54	268
Adjusted R^2	0.058	0.206	0.055	0.067	0.253	0.066

*p<0.1; **p<0.05; ***p<0.01

Notes: (1) Means and standard deviations of price adjustment factors (PAFs) are weighted by bid quantity; (2) Volatility Index (OVX) is the closing price of the CBOE Crude Oil Volatility Index on the auction date; (3) Observations are weighted by uncertainty (Days Until Delivery in (1)-(3), Volatility Index (OVX) in (4)-(6); (4) Table reports unclustered heteroskedasticity-robust standard errors.

Table 5: Test of Independent Private Values (Hickman et al., 2021)

	(1)	(2)
Mean Rival PAF	0.019	−0.056
	(0.087)	(0.090)
Sweet Crude	0.180	−0.565
	(0.715)	(0.467)
Vessel DLI	−1.876***	−1.660***
	(0.682)	(0.439)
Constant	−2.001**	−1.553***
	(0.779)	(0.520)
Polynomial in N	Cubic	Cubic
Restricted Sample	Yes	No
Observations	144	217
R^2	0.087	0.088
F Statistic	2.170** (df = 6; 137)	3.383*** (df = 6; 210)

*p<0.1; **p<0.05; ***p<0.01

Notes: (1) Restricted sample excludes auctions with any bids stipulating high minimum quantities; (2) Table reports unclustered heteroskedasticity-robust standard errors.

Table 6: Model Validation

	Sample		Model	
	Mean	Std	Mean	Std
Mean Price (PAF)	-0.69	1.71	-0.69	1.71
Quantity Dispersion	777.51	626.08	829.03	185.05
Bid-to-Cover Ratio	3.04	1.26	3.06	2.11

Notes: (1) Sample consists of 45 exclusive tranches having 3 or more bidders. Model is the calibrated uniform price basis auction. For "Sample" sample means are reported; for "Model" tranche-level expected values. (2) Dispersion defined as the standard deviation in purchase quantities within each auction. (3) Bid-to-cover ratio defined as the ratio of aggregate demand at the minimum observed price bid to quantity available.

Figure 3: Information Structures in the Vives (2011) Model

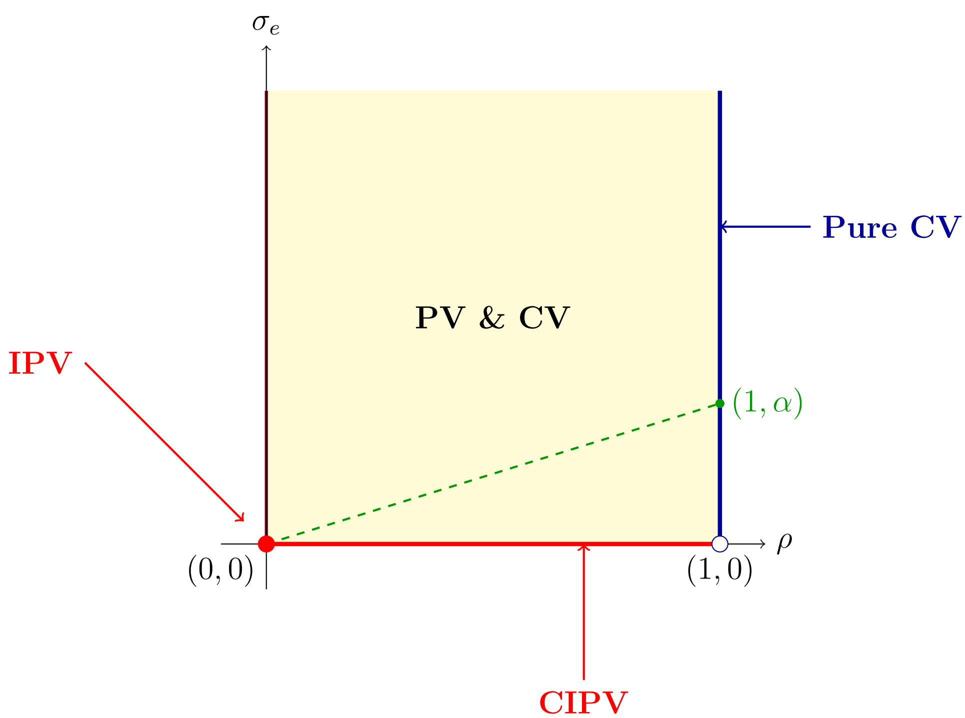

Figure 4: Simulated Rejection Rates for $(\rho, \sigma_e) = (t, \alpha t)$

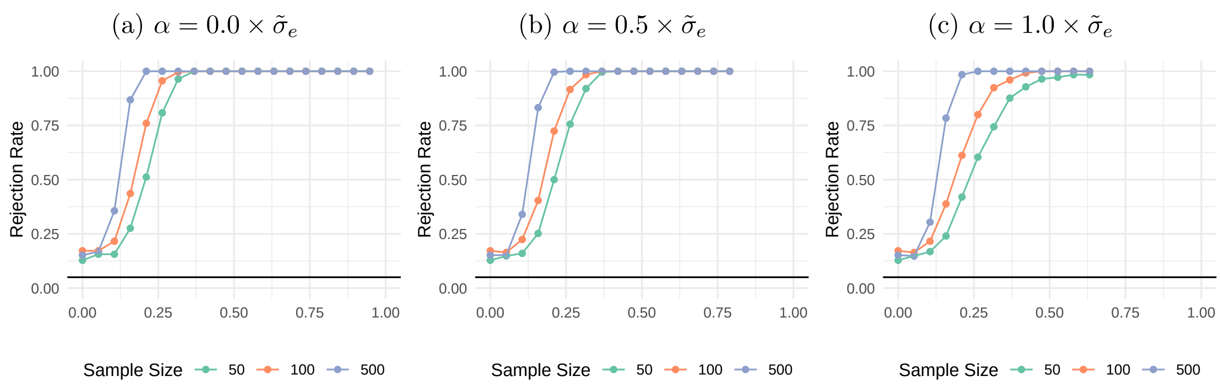

Notes: (1) As $t$increases (x-axis), the information structure is further from IPV. See Appendix B for details. (2) Approximate rejection rates from 250 simulations.