DISCUSSION PAPER SERIES

DP17673

UNDERBIDDING FOR OIL AND GAS TRACTS

Julien Martin, Martin Pesendorfer and Jack Shannon

INDUSTRIAL ORGANIZATION

ISSN 0265-8003

UNDERBIDDING FOR OIL AND GAS TRACTS

Julien Martin, Martin Pesendorfer and Jack Shannon

Discussion Paper DP17673
Published 15 November 2022
Submitted 09 November 2022

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

UNDERBIDDING FOR OIL AND GAS TRACTS

Abstract

Common values auction models, where bidder decisions depend on noisy signals of common values, provide predictions about Bayesian Nash equilibrium (BNE) outcomes. In settings where these common values can be estimated, these predictions can be tested. We propose a series of tests, robust to assumptions about the signal structure, to determine whether the observed data could have been generated by a Bayesian Nash equilibrium. In the setting of oil and gas lease auctions in New Mexico, we find evidence that participation decisions are correlated and that participants systematically underbid in light of ex post outcomes.

JEL Classification: N/A

Keywords: N/A

Julien Martin - julienmartin253@gmail.com
Nera Economic Consulting, London

Martin Pesendorfer - m.pesendorfer@lse.ac.uk
London School Of Economics and CEPR

Jack Shannon - j.b.shannon@lse.ac.uk
London School Of Economics

Underbidding for Oil and Gas Tracts

Julien Martin[*], Martin Pesendorfer[†], and Jack Shannon[‡][§]

November 9, 2022

Abstract

Common values auction models, where bidder decisions depend on noisy signals of common values, provide predictions about Bayesian Nash equilibrium (BNE) outcomes. In settings where these common values can be estimated, these predictions can be tested. We propose a series of tests, robust to assumptions about the signal structure, to determine whether the observed data could have been generated by a Bayesian Nash equilibrium. In the setting of oil and gas lease auctions in New Mexico, we find evidence that participation decisions are correlated and that participants systematically underbid in light of ex post outcomes.

Keywords: Testing, Collusion, Auctions

JEL Codes: D44, L10, L40

[*]Julien Martin: NERA Economic Consulting, Marble Arch House, 66 Seymour St, London W1H 5BT. Email: julienmartin253@gmail.com.

[†]Martin Pesendorfer: Department of Economics, London School of Economics and Political Science, Houghton Street, London UK WC2A 2AE. Email: m.pesendorfer@lse.ac.uk.

[‡]Jack Shannon: Department of Economics, London School of Economics and Political Science, Houghton Street, London UK WC2A 2AE. Email: j.b.shannon@lse.ac.uk.

[§]We would like to thank NMSLO and Vivek Bhattacharya for making the data available. Fraser Clark and Ziwen Lei provided excellent research assistance.

1 Introduction

Oil and gas production in the state of New Mexico generates a tremendous amount of revenue. Home to one of the most productive oil basins in the US, New Mexico received $4.1 billion in oil and gas tax revenue in the fiscal year 2022.[1] Rights to drill on state lands are auctioned off each month by the New Mexico State Land Office (NMSLO), but the prices paid in these auctions are difficult to reconcile with the massive value provided to leaseholders. Though only one of eight leases are drilled, we find that the average profit from obtaining a lease is still well over eight times the price paid in the lease sale. Such a dramatic difference between price paid and value is cause to suspect collusive behavior on the part of bidders in these sales. This paper examines features of the New Mexico oil and gas lease sales that would facilitate collusion and proposes a series of statistical tests to rule out the possibility of BNE bidding under a wide range of equilibria.

Methods to detect collusion in first-price auctions (e.g., Porter and Zona (1993, 1999); Bajari and Ye (2003); Chassang and Ortner (2019); Chassang et al. (2022); Kawai et al. (2022); Kawai and Nakabayashi (2022)) focus on detecting bidding anomalies inconsistent with BNE bidding. These methods provide antitrust authorities with a set of statistical screening devices which are aimed at settings where collusive bid patterns are not sufficiently sophisticated to disguise their intentions. We complement the existing methods by incorporating ex post returns to construct direct tests for BNE bidding and minimum revenue. Information on ex post returns is available in many auction settings. Ex post returns are naturally available in oil and gas lease sales as lease values are measurable using publicly available production records, but can also be used more broadly in settings in which resale or secondary markets can be used to assess object value.

We consider the sale of oil and gas leases in first-price, sealed bid auctions by the NMSLO over a twenty-two-year period from 1994 to 2015. For each auction we observe each bid

[1] Reported by the New Mexico Taxation and Revenue Department: https://www.tax.newmexico.gov/all-nm-taxes/oil-natural-gas-mineral-extraction-taxes/

submitted, along with the identity of the bidder, as well as a land survey description of the tract. By connecting publicly available data on oil and gas production to geographic descriptions of leased tracts, we construct estimates of the profit generated by each lease. Examining the winners across the sample period, we find that the bidding market is highly concentrated, with the largest four bidders holding a market share of more than 50%. Since lease sales are held in person each month, these dominant bidders had ample experience with each other and many opportunities to interact.

First, we find that bidder participation decisions are correlated across bidder pairs, conditional on ex post tract value. We can reject the null of zero correlation in participation decisions for approximately one-third of bidder pairs conditional on a positive return and one-half of bidder pairs conditional on a non-positive return. We extend this test using a biprobit model to control for ex post value and bidder-sale fixed effects and find a similar result, rejecting the null hypothesis of zero correlation at the 5% significance level for 21 of 48 bidder pairs. Since the conditional independence of bid distributions is an implication of the common values auction with independent conditional signal distributions, we reject the null hypothesis that the observed bids were generated by this BNE. We relax the assumption of independent signal distributions for our subsequent tests.

Second, we utilize the return data to test for the existence of profitable strategic deviations. For strategies to constitute an equilibrium they must maximize ex ante expected payoffs regardless of the information available to bidders at the time of the auction. We propose an underbidding test based on the Nash equilibrium condition that unilateral deviations cannot be profitable. Our test builds on the bid-scaling (winner's curse) test proposed in Hendricks et al. (1987) and is robust to the information structure available to bidders. We find that auction participants substantially underbid relative to the maximally profitable unilateral deviation. We find that when all bids of a bidder are multiplied by a factor of 3.2 holding rival bids constant, then the expected bidder payoff doubles, which is a violation of BNE bidding.

Finally, we compute the theoretical lower bound on expected auction revenue derived by Bergemann et al. (2017) and find that the average winning bid is considerably below this bound. Winning bids in a first-price auction are bounded away from zero; if rival bids are too low, upward deviation strategies (such as the ones considered in our previous test) will be profitable because the value of winning outweighs the cost of raising one's own bid. Bergemann et al. (2017) show that the lowest distribution of bids that can sustain a BNE corresponds to a particular "worst-case" equilibrium where the information structure is such that it minimizes auctioneer revenue. We compute the distribution of winning bids under this "worst-case" equilibrium using data on the number of bidders and the distribution of ex post values and compare it to the observed winning bid distribution. We find that auctioneer revenues could be more than tripled by moving from the status quo to even the worst possible BNE.

While the test results we obtain are suggestive of the presence of collusion, it is important to note that none of the tests we propose *prove* its existence. Section 8 discusses practical steps the auctioneer can take when facing bidders who seem likely to conspire to rig their bids.

Related literature

Our paper contributes primarily to the literature on statistical cartel detection in first-price auctions, see Porter (2005) for a survey. Porter and Zona (1993, 1999) find evidence that cartel bids are statistically different from non-cartel firms. Bajari and Ye (2003) devise an exchangeability test, which stipulates that under competition a bidder's own private information should be the only determinant of bids, while the identities of rival firms should not matter. Schurter (2017) and Chassang and Ortner (2019) propose tests for collusion using exogenous variations in reserve price levels. Kawai and Nakabayashi (2022) examine multi-round bidding and devise a test for correlation between the initial and subsequent bid. Kawai et al. (2022) propose a regression discontinuity approach by comparing marginal winning and marginal losing bids using bidders' incumbency or backlog status. Chassang et al. (2022) propose a test for missing mass in the bid distribution close to marginal losing and winning bids. We applied this test but could not find evidence of missing mass close to losing bids in our data.

Tests for collusion have also been applied for other auction settings. Conley and Decarolis (2016) propose tests to examine coordination in average-price auctions. Kaplan et al. (2016) provide tests for partial cartels in English auctions. Internal workings of cartels have been studied by Pesendorfer (2000) and Asker (2010). Caoui (2022) estimates damages from bid rigging.

Incentives to collude are studied more broadly in the face of antitrust authorities in Harrington (2005) and measured in the context of mergers in Miller et al. (2021) and Igami and Sugaya (2021). These papers exploit predictions obtained by the theory of repeated games (Green and Porter (1984); Rotemberg and Saloner (1986)).

Our paper also contributes to the literature on oil and gas lease auctions, see Porter (1995) for a survey on offshore auctions. Hendricks et al. (2003) test for winner's curse effects and confirm that bidders bid rationally in offshore sales. Kellogg (2014) studies the effect of price expectations on the decision to drill onshore. Kong (2020, 2021), Bhattacharya et al. (2022) and Ordin (2019) study the same NMSLO data as we do. Kong (2020, 2021) studies the relationship between first-price and English auctions. Bhattacharya et al. (2022) endogenize the drilling decision and study the optimal royalty rate design in contingent payment auctions. Ordin (2019) studies the role of tax policies in oil and gas lease sales. Hodgson (2019) studies information externalities in UK offshore drilling decisions. Kong et al. (2022) study a two-dimensional bidding system in Louisiana where bidders submit a bonus bid and a royalty rate.

The paper is organized as follows: Section 2 describes our framework. We describe the auction model and discuss the assumptions. Section 3 devises statistical testing procedures aimed at detecting collusion. Section 4 describes the market and highlights features that may facilitate collusion. Section 5 argues that bidders coordinate their bidding strategies.

Section 6 shows that bidders underbid. Bids are too low to maximize ex ante bidder profit. Section 7 estimates the lower revenue bound and shows that observed revenues fall short by a third of the bound induced by competitive bidding. Section 8 concludes.

2 Framework

Our framework is the pure common values mineral rights model as described in Bergemann et al. (2017), which contains the classic mineral rights model proposed in Wilson (1977) as a special case.

A seller has one tract for sale. Bidders $i = 1, \ldots, N$ are risk-neutral and bid for the tract. The tract has a common value v contained in a compact interval $V = [\underline{v}, \overline{v}] \subset \mathbb{R}$. The value v is drawn from the cumulative distribution function (cdf) G with support V. The value distribution is common knowledge among bidders. Bidder i additionally receives private information about the value beyond knowing the prior distribution. This information comes from a signal $x_i \in [\underline{x}, \overline{x}] \subset \mathbb{R}$ that is correlated with the value v. We denote $\mathbf{X} = (X_1, \ldots, X_N)$ the random variables and $\mathbf{x} = (x_1, \ldots, x_N)$ the realizations. The joint distribution of signals and ex post tract value is $F(\mathbf{x}, v)$. The seller announces a minimum bid, or reserve price, $r \in \mathbb{R}_+$.

Denote the set of high bidders with $W(\mathbf{b}) = \{i \quad | b_i \geq b_j, \text{ for all } j = 1, .., N \text{ and } b_i \geq r\}$, where $\mathbf{b} = (b_1, \ldots, b_N) \in B^N = [0, \overline{v}]^N$ denotes the vector of bids. Let the probability that bidder i receives the good be $q_i(\mathbf{b}) = 1/|W(\mathbf{b})|$ if bidder i is among the high bidders, and $= 0$ otherwise.

A bidding strategy for player i is a mapping $\beta_i : [\underline{x}, \overline{x}] \longrightarrow B$ from signals to bids. Let Σ_i denote the set of strategies for bidder i and let $\beta \in \mathbf{\Sigma} = \times_{i=1}^{N} \Sigma_i$ denote a strategy profile.

Bidder i's ex ante payoff from the first-price auction is given by

$$U_i(\beta) = \int_{v \in V} \int_{\mathbf{x} \in [\underline{x}, \overline{x}]^N} [v - \beta_i(x_i)] \, q_i(\beta(\mathbf{x})) \, dF(d\mathbf{x}, dv). \tag{1}$$

The profile β is a Bayesian Nash Equilibrium (BNE) if and only if $U_i(\beta) \geq U_i(\beta_i^{'}, \beta_{-i})$ for all $\beta_i^{'} \in \Sigma_i$.

Discussion of the assumptions.

Wilson (1977) and most of the subsequent empirical literature on common value auctions require stronger assumptions than stated above. Wilson assumes that the signal X_i is iid with continuous conditional cdf $F(.|v)$. The joint distribution of signals and ex post tract value is then $F(\mathbf{x}, v) = \prod_i F(x_i|v)G(v)$.[2]

Our data do not include information on bidders' signals. Proposition 4 in Laffont and Vuong (1996) establishes that the signal distribution $F(.|v)$ cannot be identified from bid data and ex post values alone. For example, monotone rescaling of signals results in observationally equivalent signal distributions. See also Somaini (2020) for identification results with interdependent signals. While signals are not identified, we can explore statistical properties of signals by examining bids instead. Bids are observed and are a strict monotone function of signals. If signals are independently distributed, then bids must be as well. We shall consider tests of this assumption using bids instead of signals in Section 5.

Prior empirical work on common value auctions typically adopts the Wilson BNE and imposes additional assumptions on the information structure to guarantee identification from observables. Bhattacharya et al. (2022) assume that bidder receive noisy signals of the quantity of oil in a tract. Hendricks et al. (2003) assume that the signal is an unbiased estimate of the ex post return conditional on winning. We shall consider tests of the null hypothesis that bids satisfy the conditional independence assumption in Section 5. A rejection of these tests may be indicative of bidder coordination prior to the auction, but could also suggest

[2]Based on the iid signal assumption Wilson (1977) and Milgrom and Weber (1982) characterize the symmetric BNE. Let $Y_i = \max_{j \neq i} X_j$ and $F_{Y_i|X_i}(.|.)$ be the conditional distribution of Y_i given X_i. Let $u(x_i) = E[v|X_i = x_i, Y_i = x_i]$ be the expected value conditional on the own signal being x_i and the high rival signal being at most x_i. The equilibrium strategy satisfies the first-order differential equation $b'(x) = [u(x) - b(x)] \cdot \frac{f_{Y_i|X_i}(x|x)}{F_{Y_i|X_i}(x|x)}$ with boundary condition $b(\underline{x}) = u(\underline{x})$. The solution is

$$\beta_i(x_i) = u(x_i) - \int_{\underline{x}}^{x_i} L(y|x_i)du(y) \quad \text{where} \quad L(y|x_i) = \exp[- \int_y^{x_i} \frac{f(x|x)}{F(x|x)} dx.$$

that the iid signal assumption of Wilson is not satisfied (if, e.g., signals are correlated with elements other than the ex post return). With this caveat in mind, our main statistical analysis will be based on the weaker set of assumptions described above which are robust to alternative information structures.

To summarize, our statistical analysis departs from the prior literature by using a weaker set of assumptions that is robust to all information structures, including the one proposed by Wilson. Our approach is robust to the specification of the signals and details of the Bayesian Nash equilibrium.

3 Testable Implications

This section describes testable implications of BNE bidding. We will formulate suitable statistical tests of these implications using the publicly available data on oil and gas lease sales from the New Mexico State Land Office (NMSLO).

The sealed-bid first-price auction has bidders submitting sealed bids and awards the item to the high bidder at his bid price. The identities of potential bidders are publicly known before every auction. On the day of the auction, the sealed bids are publicly revealed, and the high bidder wins. We let $\mathbf{b}^t$ denote the vector $(b_1^t, ..., b_N^t)$ and adopt the convention that $b_i^t = 0$ when potential bidder i refrained from bidding for lot t. We denote with z^t any information that is publicly available at time t, such as the oil and gas spot (and future) prices, that may affect bidders' signals $\mathbf{x}^t$ and thus bids $\mathbf{b}^t$. The variable v^t denotes the ex post return, which we calculate from the publicly observed drilling and production data.

We make the following assumption on the data generating process (DGP).

Assumption 1. *The observed data are* $(\mathbf{b}^t, v^t, z^t)_{t=1}^T$ *where* $(\mathbf{b}^t, v^t)$ are *identically and independently distributed conditional on exogenous covariates* $z^t \in Z$*.*

The assumption is commonly imposed in market games, see Tamer (2003).

We shall focus on three central implications of the mineral rights model, each requiring a decreasing amount of structure. First, we consider the classic Wilson model in which the submitted bids are independently distributed conditional on ex post returns and publicly available information at the time of the auction. Second, we relax the independence assumption and examine the null hypothesis that bidding strategies maximize ex ante expected returns, that is, that no bidder can systematically deviate from the equilibrium and receive strictly larger expected profits. Third, the distribution of winning bids must satisfy the robust lower revenue bound described in Bergemann et al. (2017). The last two properties hold regardless of bidders' information and the Bayesian Nash equilibrium. We shall describe these hypotheses in turn.

3.1 Independence

The data include detailed information on ex post drilling outcomes which allow us to calculate ex post returns for bidders, which we use as a control variable. Evidence of correlation in bids conditional on ex post returns is indicative of pre-play communication, which would be a violation of the Bayesian Nash equilibrium condition in Wilson's mineral rights model.

Implication 1. *Consider the assumption of Wilson (1977). The bids (signals) B_i and B_j are independently distributed for all $i, j \in N$ conditional on the ex post value realization v.*

The null hypothesis of conditional independence is

$$H_0^{B|X} : B_i \perp B_j | v \text{ for all } i, j \in N, \tag{2}$$

with the alternative its negation. The null can be tested for individual bidders or for bidder pairs. A violation of the null hypothesis suggests that the data were not generated from the BNE in the Wilson model. This result could indicate that the data were not generated from a BNE (e.g. because of collusion among a subset of bidders) or that the data were generated by a BNE in a game with a different information structure (e.g. if value signals are truly correlated even after conditioning on ex post values).

3.2 Best Response Test

Next, we relax the assumption on the information structure and go beyond the Wilson model. A simple yet powerful test examines the BNE condition ex ante using equation (1).

Bidder i's payoff realization from an auction is given by

$$U_i(b, v) = [v - b_i]\, q_i(b_i, b_{-i}).$$

Consider a unilateral deviation $\phi : B \longrightarrow B$ which results in the modified payoff realization

$$U_i(b, v|\phi) = [v - \phi(b_i)]\, q_i(\phi(b_i), b_{-i}). \tag{3}$$

The BNE condition requires that a unilateral deviation cannot be profitable on average. Consider the ex ante payoff under the deviation strategy:

$$U_i(\beta|\phi) = \int_{v\in V} \int_{\mathbf{x}\in[\underline{x},\overline{x}]^N} [v - \phi\left(\beta_i(x_i)\right)]\, q_i(\phi\left(\beta_i(x_i)\right), \beta_{-\mathbf{i}}(\mathbf{x}))\, dF(d\mathbf{x}, dv). \tag{4}$$

This leads us to the following implication of BNE bidding.

Implication *2. In any BNE under any information structure and for any bidder i, the function ϕ that maximizes the ex ante expected payoff in equation (4) must be the identity mapping.*

This property can be used to detect any deviation from BNE bidding. Our goal is to detect whether bidders systematically underbid. We follow Hendricks et al. (1987) in considering linear deviations, that is, deviation strategies that multiply all bids by a scalar $\alpha > 0$ such that $\phi(b) = \alpha \cdot b$. Letting $\alpha^* = \arg\max_\alpha U_i(\beta, \alpha)$, Implication 2 of BNE bidding leads to the following null hypothesis:

$$H_0 : \alpha^* = 1.$$

Originally, this type of best response test was aimed at testing for the presence of the winner's curse (systematic overbidding, or $\alpha^* < 1$). We expand this test to check optimality of bidding more generally.

3.3 Lower Bound on Revenues

Our third test of BNE bidding utilizes the distributional bound on winning bids characterized in Bergemann et al. (2017). This bound is a useful tool for empirical analysis as it depends only on the ex post return distribution, which we can recover from the data. Prior empirical work specifies the information distribution of bidders or the details of the Bayesian equilibrium. Our approach departs from the prior literature and is robust to the specification of the information structure, signals, and other details of the Bayesian equilibrium. We test whether observed returns exceed the lower bound in expectations.

The lower bound is tight in the sense that it emerges in the bidding equilibrium with the "worst-case" information structure in the sense that expected revenues to the seller are lowest. Bergemann et al. (2017) establish that signals x_i drawn from $G(X_i)^{1/N}$ (recall that G is the cdf of the distribution of values) and correlated with the value so that $\max x_i = v$ achieve this bound.

Winning bids in such an equilibrium are given by[3]

$$\underline{\beta}(v;r) = \begin{cases} \frac{1}{G(v)^{(N-1)/N}} \left[r \cdot G(\hat{v})^{(N-1)/N} + \int_{\hat{v}}^{v} \frac{N-1}{N} \cdot \frac{x}{G(x)^{1/N}} dG(x) \right], & v > \widehat{v}; \\ 0, & v \leq \widehat{v}, \end{cases}$$

[3]It is useful to observe that the minimum bid is a solution to the differential equation

$$\frac{N-1}{N} \cdot \frac{v - \beta(v)}{G(v)} dG(v) - \beta'(v) = 0$$

with boundary condition $\beta(\hat{v}) = r$.

where the threshold value $\widehat{v}$ solves

$$\frac{\int_{\underline{v}}^{\hat{v}} x \, dG(x)}{\int_{\underline{v}}^{\hat{v}} dG(x)} = r.$$

The minimum ex ante expected revenue is

$$R_{\min}(r, G, N) = \int_{\hat{v}(r)}^{\overline{v}} \underline{\beta}(x; r) \, dG(x).$$

Implication 3. *In any BNE under any information structure, the expected revenue from the auction must exceed the minimum revenue bound $R_{min}(r, G, N)$.*

Under the null of competitive bidding, we assume that the observed bids were generated by a profile of BNE bidding strategies and denote the expected winning bid, conditional on tract value v and reserve price r, as $\beta_w(v; r)$. We test Implication 3 using the null hypothesis

$$H_0 : \int_{v \in V} \left[\beta_w(x; r) - \underline{\beta}(x; r) \right] dG(x) \geq 0. \tag{5}$$

The lower bound used in this test is robust to details of the Bayesian game being played and the information available to bidders. The bound remains valid if bidders are asymmetrically informed, completely uninformed, or completely informed about the tract values. The bound also applies when there is unobserved auction heterogeneity, that is, when bidders publicly observe a part of the common value that is not recorded in the data. The bound is also robust to risk aversion – if bidders were risk averse, we would expect more aggressive bidding, pushing revenues even higher above the bound.

Next, we shall describe the market and data.

4 Data

We study oil and gas lease sales held monthly by the New Mexico State Land Office (NMSLO) on the third Tuesday of each month between 1994 until 2015.[4] Every month the NMSLO distributes a list of leases which are sold at auction next month. Bidders can nominate an area for auction. Leases typically cover a land area of 320 acres.

Lease Sale Procedure

It is a legal requirement that the NMSLO awards leases by one of two types of auction formats[5]: (i) sealed-bid first-price auction and (ii) open-outcry English auction. For most of the sample period both formats were used in each monthly sale.[6] The assignment of auction format is mostly random, although conversations with the leading auctioneer during the sample period suggest that, in the case of split tracts (a tract larger than 320 acres that is split into two separate lots for the sale), the first price auction is used for the larger of the two halves or for the part closest to an existing lease owned by the bidder that nominated the tract. The monthly lease sale proceeds in two stages: First, the sealed bids for the first set of leases are unsealed and announced publicly. Every lease is awarded to the high bidder. Second, the set of English auction leases are awarded in sequence by means of an open outcry auction where all eligible bidders are present in the room. Every month the same set of bidders interact. We do not include auctions taking place after 2015 in our sample as the post-award production record is incomplete.

The lease duration is five years, during which time the winning bidder can drill a well. If oil or gas is found, then the lease can be extended until all minerals have been extracted. Any revenues arising from the well are subject to various revenue taxes and royalty payments at a rate depending on the type of lease. Additionally, leaseholders are charged a negligible

[4]Data from NMSLO have been studied in the prior literature. Kong (2020, 2021) studies the relationship between first-price and English auctions. Bhattacharya et al. (2022) study the effect of post auction drilling decisions on the optimal auction design in terms of the royalty rate.

[5]The legal setting is described in New Mexico Statues, Chapter 19, Article 10, Section 19-10-17: https://law.justia.com/codes/new-mexico/2019/chapter-19/article-10/section-19-10-17/

[6]During the summer of 2016 the auctions moved online. Starting in 2019 most of the sales were conducted in first-price format only.

rental rate, typically \$0.50 or \$1.00 per acre. The winning bidder pays corporate taxes on profits, which is a proportional deduction and thus neutral to the bidding calculus. The winning bidder pays the bid bonus immediately after the auction and receives the net return on future and uncertain minerals extracted.[7]

Our data contain detailed information on ex post drilling outcomes and production until 2021 which enables us to calculate the future return following the prior literature. After the auction, the winner of the tract obtains the tract lease and the subsequent oil and gas extraction is publicly observed. Bhattacharya et al. (2022) and Hendricks et al. (2003) construct the value measure using the realized value of minerals extracted from ex post drilling activity. We follow this definition of the common value and construct our value estimates by matching publicly available production data to leases. We obtain a list of all oil and gas wells in New Mexico from the New Mexico Oil Conservation Division (OCD) describing the location of each well. We then match these well locations to the geographic descriptions of leased tracts provided by the NMSLO in the letting announcements. For each lease we aggregate the monthly oil and gas production of each well (collected from monthly production reports submitted to the OCD) and weigh them with deflated crude oil and gas prices. To account for production delays, we discount all returns to the date of the auction using a five percent annual interest rate. We account for the royalty rate, which varies by the type of lease, and deduct a revenue tax of 7.1%[8]. Our final gross revenue measure equals the realization of the discounted ex post value of the oil field net of royalties and taxes. Our

[7]If during future production it is found that the well can be used to extract minerals for multiple leases, then bidders are by law required to enter a pooling agreement. While the share of leases with pooling agreements is relatively small in our data (less than four percent), it forces bidders to engage and cooperate with each other on those leases.

[8]According to Chapter 2 of the "Decision-Makers Field Guide (2002)," (available online at https://geoinfo.nmt.edu/publications/guides/decisionmakers/2002/dmfg2002_complete.pdf) there are six taxes imposed directly on oil and gas extraction and processing: (i) severance tax which amounts to 1.875% during the first 5-7 years of production and then increases to 3.75%; (ii) conservation tax of about 0.19%; (iii) emergency school tax of 3.15% for oil and 4% for gas minus drilling credit which is given some times; (iv) ad valorem production tax of about 0.39%; (v) natural gas processors tax of 0.45% and (vi) ad valorem equipment tax of 0.07%. Totaling these taxes amounts to about 7.1%. The rate of 7.1% is also reported as tax revenues obtained in the year 2000 on the value of oil and gas reported by the Taxation and Revenue Department, see Figure 5 in Decision-Makers Field Guide (2002). Following Ordin (2019) we observe that corporate profit taxes do not affect the bidding calculus as the tax rate applies proportionally.

gross revenue measure is a lower bound on revenues, as it is based on observed production and does not include potential future production beyond 2021[9]. We define the common value v as the net return, which equals the gross revenue measure minus the well cost.

Well costs are measured following the formula described in Kellogg (2014), which is based on drilling rig rental costs predicted by oil future prices. Kellogg studies oil extraction in the Texas region of the Permian Basin, while most of the wells we study fall in the New Mexico region of the Permian Basin. Kellogg notes a significant positive correlation between 18-month-ahead oil prices and rig dayrates. He regresses daily drilling rig rental rates on 18-month-ahead oil prices and obtains an R-squared of 0.64. This regression is used to predict the daily rig rental rate, which is then multiplied by the expected number of drilling days to get the expected rental cost, which is in turn multiplied by three to produce total expected drilling costs.[10] We update Kellogg's estimate of 19.2 days to drill a well on average with a more conservative 27.4 days, which emerged from a sample of 59 New Mexico drilling cost records on pooling agreements which are publicly available[11]. Our well cost estimates are about 40% higher than Kellogg's.

To check the accuracy of our well cost estimates, we benchmark our cost estimates against two data sources: (i) a sample of 313 publicly available New Mexico well cost estimates obtained from pooling agreements and (ii) well cost estimates by the US Energy Information Administration[12]. Our well cost estimates are marginally higher than the New Mexico benchmark sample and in the range of the US Energy Information estimates. Our well cost

[9]An alternative measure of gross revenues uses the oil and gas future prices at the time of the auction as weights to discounted quantities instead of the realization of the oil and gas prices at the time when production takes place. The resulting gross revenue measure is very similar in magnitude but on average slightly larger than the measure obtained using the realized prices.

[10]The scaling factor of three emerged from conversations with industry members who estimated that rig rental costs constitute on average one third of total drilling costs.

[11]The New Mexico well cost sample is obtained from publicly available reports on pooling agreements. Parties engaging in a pooling agreement are required by law to submit an Authorization for Expenditure (AFE) that describes the anticipated cost of a proposed well to New Mexico's Oil Conservation Division. We parse AFEs filed between 2008 and 2021 and extract reported number of days drilled for 59 records. The OCD website is https://wwwapps.emnrd.nm.gov/OCD/OCDPermitting/Data/Hearings/Cases.aspx

[12]We parse AFEs filed between 2008 and 2021 and extract total drilling and completion costs from 313 records. The US Energy Information Administration reports well cost estimates in https://www.eia.gov/analysis/studies/drilling/pdf/upstream.pdf.

Table 1: Summary Statistics for Awarded Tracts

	All	First-Price	English
Number of Auctions	9,717	4,535	5,182
Gross Revenue (minus royalty and tax payment)	626.1	702.1	559.7
	(3,817)	(4,362)	(3,265)
Well Cost	192.8	195.6	190.3
	(755.2)	(794.9)	(718.7)
Net Revenue v	433.3	506.4	369.4
	(3,229)	(3,711)	(2,737)
Winning Bid	52.82	58.68	47.70
	(113.5)	(138.5)	(85.0)
Reserve Price	4.71	4.53	4.87
	(4,298)	(4,302)	(4,474)
Fraction Drilled	0.125	0.119	0.131
	(0.331)	(0.323)	(0.337)
Well Cost of Drilled Tracts	1,541	1,649	1,455
	(1,576)	(1,713)	(1,453)
Fraction Productive	0.110	0.106	0.114
	(0.313)	(0.308)	(0.317)
Disc Revenue of Productive Tracts	5,681	6,606	4,924
	(10,175)	(11,842)	(8,509)

The data consist of all awarded auctions between 1994 and 2015. Dollar figures are measured in thousand of 2000 US dollars. Standard deviations are in parentheses.

estimates do not take heterogeneity in well costs across bidders into account. We expect drilling costs differences across leases to be small, owing to the relatively uniform geology of the Permian Basin. While cost differences may exist across bidders, we follow Hendricks et al. (2003) and Bhattacharya et al. (2022) in assuming that these differences are small in magnitude or not known at the time of the auction.

Descriptive Summary Statistics

Table 1 shows that 4,535 sales were held using first-price auctions while 5,182 using an English auction format. Table 1 considers only pre-2016 sales as the initial drilling decision can occur at the end of the lease term, as shown in Bhattacharya et al. (2022). All dollar magnitudes are deflated using 2000 dollars.

Strikingly, the bonus bid is very small relative to tract value. The bonus bid equals \$53,000 on average, which amounts to 12 percent of the average tract value. In comparison, the offshore lease sale literature has shown bids being much closer to the value of the tract. Hendricks et al. (1987) report winning bids in offshore sales equal 76 percent of tract value for wildcat sales and 49 percent for drainage sales.

A second surprising element is that only a small fraction of awarded tracts are drilled: about 12.5 percent.[13] Most drilled wells are productive (i.e. the well produced oil or gas). The drilling rate and productivity rate are both low in comparison to that of offshore tracts, which Porter (1995) finds to be 78 percent and 35 percent, respectively. We hypothesize that the low initial drilling rate is a result of lease hoarding, which appears common for onshore leases[14]. We shall provide further evidence on the number of undrilled leases hoarded by individual bidders below.

Interestingly, the auctioneer's revenues are higher for first-price auctions than for English auctions, both in terms of royalty payments and bonus bids, see also Kong (2020). For split tracts, in which both auction formats were used and the assignment (according to conversations with the lead auctioneer) is essentially random, the first-price auction generates 23% higher cash bonus bid and also a 37% higher royalty return per acre. We can reject the null of identical bonus bids across the two formats at the one percent significance level. Yet, the null of identical royalty returns cannot be rejected. During the year 2019 the NMSLO began awarding leases exclusively in the more favorable first-price format.

The revenue ranking is surprising in the light of the classic theoretical work on symmetric BNE bidding equilibria in standard auction formats. Milgrom and Weber (1982) derive that English auctions generate more revenues than first-price auctions on average. One explanation for the revenue superiority of the first-price auction format is that some coordination or collusion arises in English auctions. Avery (1998) shows that bidders may use initial jump bids to signal their intention to rivals, which gives rise to multiple equilibria in English

[13]The drilling rate estimate is almost identical to the onshore drilling rate reported in the prior literature, see Bhattacharya et al. (2022).

[14]Accoding to a Wilderness Society's article from December 15, 2015, hoarding is common in the oild and gas industry. For instance, suspension of federal leases has affected 3.25 million acres in April 2015. See https://www.wilderness.org/articles/blog/land-hoarders-oil-and-gas-companies-are-stockpiling-your-public-lands.

auctions, some of which may have drastically reduced revenues. Indeed, we observe jump bidding in English auctions conducted online, where the timestamp of each bid is recorded.[15]

Suspiciously low English auction bids arise also in split tract sales, where bidders should arguably have the same value estimate for both halves. There are 335 occasions where a bidder failed to win the English auction although the bidder submitted a bid in the first-price auction that was (substantially) higher in per-acre terms than the selling price in the English auction. On 194 of these occasions, the bidder failed to win both the first-price auction and the English auction. On average the bidder's losing first-price auction bid was 120 percentage points higher than the final English auction price. While these bid patterns seem odd, they can in fact arise as a BNE when bidders have beliefs that they will be outbid in the English auction. Since BNE bidding in English auctions may resemble coordination or collusion, our subsequent analysis focuses on first-price auction sales.

Factors Facilitating Collusion

There are several factors in the lease sale market that may facilitate coordination or collusion. We think of collusion as an implicit or explicit arrangement to limit competition among market participants and to increase profits.

The market we study is concentrated, with three firms winning half of all leases sold at auction. A small set of firms coordinating their actions can have a big impact on market price. Table 2 reports summary statistics for bidders who won more than 100 leases between 1994 and 2021 and together account for two thirds of all bids submitted and 60% of auction awards. The table also includes a "fringe" bidder accounting for all remaining bids. We report dollar measures on average across all auctions won by the bidder. Yates Petroleum Corp has a market share of about 38% in the number of leases with more than 1,000 active (undrilled) leases (320,000 acres) held during any calendar year between 2000 and 2015. It

[15]The NMSLO online English format is similiar to eBay sales where the current standing winning price is revealed to rival bidders and not the submitted bid. In the online English lease sale in January 2019, Slash Exploration LP started the bidding with a bid substantially above the reserve, and two attempts by rival bidders during the next sixteen hours to outbid Slash Exploration failed, resulting in Slash Exploration winning the lease.

Table 2: Bidding Returns for Top Bidders

Bidder	No of Bids	No of Wins	Return v	Bid b	ROI
			on average across auctions won		
YATES PETROLEUM CORP	5,810	4087	325.26	30.67	18.90
DANIEL E GONZALES	828	592	979.19	65.57	11.78
DOUG J SCHUTZ	838	587	353.30	64.21	2.22
CHASE OIL CORPORATION	348	249	50.79	39.96	0.21
FEDERAL ABSTRACT COMP.	381	205	312.75	12.87	2.13
SLASH EXPLORATION LP	683	164	11.31	6.33	5.95
FEATHERSTONE DEV. C.	376	149	326.37	22.40	6.44
MARBOB ENERGY CORP	278	130	1,271.33	78.40	12.85
BAR CANE INC	220	129	601.11	60.95	4.12
RONALD MILES	369	122	1,148.06	91.09	8.13
THE BLANCO COMP.	617	103	602.69	14.39	90.94
FRINGE	N/A	4280	390.39	61.44	10.11

The data consist of all awarded auctions between 1994 and 2021. Dollar figures are measured in thousand 2000 US dollars.

was acquired by EOG Resources in 2016 for $2.5 billion. Yates Petroleum operated beyond the New Mexico region and held about 1.5 million acres in at least seven US states at the time of acquisition. The next largest bidders are Daniel E Gonzales and Doug J Schutz, both of whom have a 5% market share each, which amounts to more than 500 leases. These two bidders held on average more than 130 active (undrilled) leases during the period 2000 to 2015.

The rate of return from winning an auction (ROI), measured by the profits (the tract value minus the bonus bid) divided by the bonus bid, is substantial for all bidders, equaling more than 1,400 percentage points on average across bidders. The high percentage arises as the lease acquisition cost is very low relative to the return. We shall examine the null that the bonus bid is too low to be consistent with competitive bidding in Sections 6 and 7.

Leases are homogenous products that can be resold in the future. Competition is only in terms of price, so a cartel need only coordinate in the price dimension to collude. Sales occur regularly at monthly intervals, with bidders gathering in person for each sale. Consequently, bidders know the identity of other potential bidders before they bid. Bidders may have formed relationships with each other at prior sales or as a result of pooling agreements they are required by law to enter into when a well spans multiple leases. Bidders participating in the NMSLO's auctions have faced allegations of illicit conduct in other states.[16] Multiple leases are awarded at each sale date, allowing bidders to divide the market without using side payments. Additionally, leases can be resold at subsequent periods, providing a mechanism for bidders to implement a suitable market division. The frequency of sales makes it costly for firms to deviate from any agreement. To summarize, the market exhibits characteristics that facilitate collusion. It is a concentrated market, a homogeneous product is sold, multiple leases are sold at every sale date, and sales occur at regular monthly intervals.

Next, we shall conduct statistical tests to examine whether we can reject the null of competitive bidding.

5 Conditional Independence Test

Competitive behavior requires that bidders submit their bids independently of each other conditional on the information available to them. Bidding strategies cannot be coordinated or correlated; a player's strategy should be a function of their signal only. Coordination among competing bidders or information sharing is not legal at auctions. In contrast, when bidders coordinate or communicate prior to the auction, then we may expect bids to be correlated beyond the information available to bidders individually. This section considers tests aimed at distinguishing these two hypotheses based on the assumptions of Wilson (1977).

Wilson's mineral rights model assumes the signal x_i is drawn independently from a conditional cdf $F(x_i|v)$, where v denotes the common value of the oil field . We do not observe

[16] On March 15, 2012, the US Department of Justice filed a law suit alleging bid-rigging in Colerado, see https://www.justice.gov/atr/case/us-v-sg-interests-i-ltd-et-al. Reuters reported on June 25, 2012, that email exchanges between Chesapeake Energy Corp and a competitor apparently intended to avoid bidding against each other in Michigan, see https://www.reuters.com/article/us-chesapeake-land-deals-idUSBRE85O0EI20120625. On March 1, 2015, the US Department of Justice indicted the CEO of Chesapeake Energy Corp for bid rigging in Oklahoma, see https://www.bloomberg.com/news/articles/2016-03-01/chesapeake-co-founder-mcclendon-indicted-over-lease-bid-rigging.

the signal realizations $\mathbf{x}$ but do observe the bids $\mathbf{b}$ and the ex post outcome v. In Wilson's BNE a bid is a strict monotone function of a bidder's signal, $b_i = b(x_i)$. Since signals are drawn independently from $F(x_i|v)$, and bids are a function of one signal only, bids will be distributed independently after conditioning on ex post returns. These assumptions lead to the following null hypothesis for competitive bidding:

$$H_0^{B|X} : b_i \perp b_j | v \text{ for all } i, j \in N, \tag{6}$$

with the alternative hypothesis:

$$H_1^{B|X} : b_i \not\perp b_j | v \text{ for some } i, j \in N.$$

As most bids in our data are equal to zero (the convention we adopt to represent that a bid was not submitted), our test statistic aggregates bids into the binary bid submission variable $s_i = \{1 \text{ if } b_i > 0; \ = 0 \text{ otherwise } \}$.

Wilson additionally assumes that the conditional cdf F is identical for all bidders. Asymmetries in the distribution of bids could arise if, e.g., the variance in the signals differs across bidders. We are not concerned with testing for the presence of such asymmetries; here we only assume that bids are independently (not necessarily identically) distributed.

Our baseline independence test examines pairwise correlation in bid submission decisions conditional on ex post returns using Pearson's chi-squared test. We partition the ex post outcome v into two sets depending on whether oil is found or not: (i) positive returns and (ii) non-positive returns. We consider separate tests for both subsets. The first subset of returns contains a range of potential outcomes, whereas the second is essentially a singleton and thus well-suited for the independence test.

To construct bidder pairs we consider the set of "regular" bidders described in Table 2. These bidders account for 65% of all bids submitted. We consider pairwise independence tests of the null, $s_i \perp s_j | v$ for any pair $i, j \in N$ where the ex post return variable v is partitioned

Table 3: Correlation in Bid Decisions Conditional on Ex Post Value

	Pairs	H_0 of Independence Rejected	Sign of Correlation Positive	Negative
Positive Return: $v > 0$	15	6	9	6
Non-Positive Return: $v \leq 0$	46	19	31	15

Test results of the null of pairwise independence in the bid submission decision are reported. The significance level is 10 percent. The data consist of all pairs in which both bidders are active on sales involving at least 100 auctions in case of positive returns, and at least 200 auctions in case of non-positive returns.

into two sets, non-positive $\{v : v \leq 0\}$ and positive $\{v : v > 0\}$. We conduct the test for every bidder pair separately. For bidders who are inactive during part of the sample period, we take sales into considerations only if both bidders where active and submitted at least one bid during the sale. Finally, we consider bidder pairs that overlapped, or interacted, on at least 100 auctions for the case of positive returns and 200 auctions for the case of non-positive returns.

Table 3 indicates that the conditional independence test is rejected for 45% of bidder pairs at the ten percent significance level. The test results conditional on non-positive returns has a larger sample size and yields 50% of test results rejecting the null of independence. To illustrate that the magnitude of the correlation coefficients is economically meaningful we select two bidder pairs with significant correlation. In the case of positive ex post revenues, bidder 4 was 2.5 times as likely to submit a bid (45.3% versus 17.6%) when bidder 3 bid as well. With non-positive revenues bidder 2 was 2.1 times as likely to submit a bid (37.7% versus 17.7%) when bidder 7 bid as well.

The test results in Table 3 suggest that bidders coordinate on the participation decision, which is inconsistent with Nash equilibrium bidding in Wilson's mineral rights model. It could be argued that such patterns of coordination could also arise due to exogenous variations over time. For example, bidder pairs may be more active in certain seasons, or respond in the same way to variation in oil prices or any other exogenous shock. For example, a pattern of bid rotation would emerge when bidders are less likely to bid if a sizable number of leases have been won in the preceding sale.

Table 4: Biprobit Correlation Coefficients.

Pairs	H_0 of Zero Correlation rejected at			Sign of Correlation Coefficient	
	10%	5%	1%	Positive	Negative
48	22	21	12	33	15

Test results are reported for the null of a zero correlation coefficient in the bivariate probit for bidder pairs. The data consist of all pairs in which both bidders are active on sales involving at least 200 auctions. Explanatory variables include ex post return, ex post return squared and bidder specific sale-date fixed effects.

To account for these alternative explanations, we exploit the timing of individual auctions. Our data have multiple auctions taking place on the same date. A sale occurs once a month, and all auctions within a sale have an identical bid submission deadline. In total our data have 259 sales dates between 1994 and 2015 with an average 18 first price auctions taking place per sale.

We augment the above independence test to additionally control for date fixed effects z interacted with bidder-fixed effects. We test the null, $s_i \perp s_j | v, z$ for any pair $i, j \in N$ with a parametric bivariate probit framework. The biprobit controls for bidder-specific date effects and ex post return values (linear and quadratic). The correlation coefficient in the biprobit measures the bidder-pair correlation not accounted for by ex post return and date fixed effects z^t.

Table 4 indicates that a quarter to one half of bidder pairs have correlation coefficients that are significantly different from zero, depending on which significance level is considered. The majority of coefficients are positive. Consequently, we can reject the null of independence in the mineral rights model. The evidence suggests that if the game played by bidders is the same as in the mineral rights model, bid submission decisions are coordinated. This coordination arises among a quarter to one half of bidder pairs mentioned in Table 2.

The evidence so far has been inconsistent with competitive bidding in the mineral rights model, but could be produced by a different information structure. In the presence of unobserved heterogeneity – characteristics of the tract that are (i) informative of its value, (ii) observed by bidders, and (iii) not recorded in the data – some bidder pairs would exhibit

Figure 1: Histogram of Biprobit Correlation Coefficient Estimates

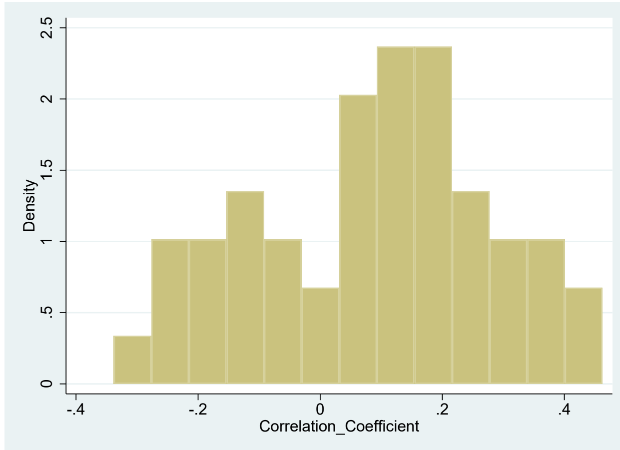

correlation in participation decisions driven by the unobserved heterogeneity. Because unobserved heterogeneity affects the value signal in the same direction (either positively or negatively) for all bidders, any correlation induced by it should be positive, see Krasnokutskaya (2011). To examine whether the correlation we observe could be explained by unobserved heterogeneity, we examine the distribution of participation correlation coefficients across bidder pairs.

Figure 1 plots the histogram of correlation coefficient estimates. Some bidder pairs appear to refrain from bidding against each other, while the majority of bidder pairs complement each other in bid submission. There are both sizable negative and positive correlation coefficients, with surprisingly little mass at zero. The histogram differs from that of a normal distribution in that it has a hump at -0.2 and another hump at 0.2. Such a bimodal distribution is inconsistent with the strictly positive correlation that would arise under unobserved heterogeneity.

The empirical evidence considered so far cannot be reconciled with competitive bidding under the mineral rights model. Having rejected the null of independent bid distributions, we are left with two possibilities. First, the bidding strategies that generated the data could be coordinated rather than competitive, e.g. some bidders refrain from bidding or submit "phony" bids. Alternatively, the correlation in bids could be explained by the underlying information structure mediating ex post returns and bids. If signals are positively correlated across some pairs and negatively correlated across others, the observed patterns of positive and negative bid correlation could arise in a competitive equilibrium. The next two sections consider statistical tests of bidding in the common values BNE that are robust to the underlying information structure.

6 Underbidding

This section examines whether bidders systematically under- or overbid.

If bidders coordinate in order to suppress bid payments, then such behavior will be detectable by finding the existence of a profitable deviation in the bidder's choice problem. Since we observe all bids, as well as the ex post return, we can measure the observed average payoff and use it to test the null that a systematic deviation cannot be profitable.

We develop a test procedure that is applicable regardless of the underlying information structure. We examine deviations from observed bidding in which all bids of a bidder are multiplied by a positive scalar, holding rival bids fixed. Of course, richer deviation strategies can be permitted and the test augmented. Nevertheless, in our case, even a scalar deviation results in a substantial profit increase.

Recall the null hypothesis

$$H_0 : \alpha^* \equiv \arg\max_{\alpha} \int_{v \in V} \int_{\mathbf{x} \in [\underline{x}, \overline{x}]^N} [v - \alpha \cdot \beta_i(x_i)] \, q_i(\alpha \cdot \beta_i(x_i), \beta_{-\mathbf{i}}(\mathbf{x})) \; F(d\mathbf{x}, dv) = 1.$$

We estimate the expectation using the sample average to compute the following test statistic:

$$\widehat{\alpha}^* = \arg\max_{\alpha} \frac{1}{T} \sum_t \left[v^t - \alpha \cdot b_i^t\right] q_i(\alpha \cdot b_i^t, b_{-i}^t),$$

where the objective function is the average payoff realization when bid b_i^t in all auctions is multiplied with the scalar parameter α.

Table 5: Best Response Test: Optimal Bid Scalar $\hat{\alpha}^*$

	Overall	Bidder	
		Top-5	Non-Top-5
Bid Scalar Estimate $\hat{\alpha}^*$	3.18	3.19	3.72
10th and 90th Quantile	[2.32,3.36]	[3.07,3.66]	[2.95,3.82]

The confidence region is estimated using 199 bootstrap samples by resampling from the set of all auctions using the stratified bootstrap.

A slope estimate $\widehat{\alpha}^* < 1$ suggests evidence of the winner's curse or risk aversion, see Matthews (1983) and Maskin and Riley (1984). On the other hand, $\widehat{\alpha}^* > 1$ suggests that bidder i underbid. Hendricks et al. (1987) propose a precursor bid scaling test and (weakly) reject the winner's curse in offshore sales. Kong (2020) develops an alternative approach that does not rely on ex post returns and finds evidence of risk aversion.

We obtain the sampling distribution for the test statistic using a stratified bootstrap to account for correlation due to the exogenous variables z^t by resampling auctions inside strata (sale dates). The approach is consistent by Assumption 1 as auction outcomes within a sale date are iid with identical exogenous variable z^t, see Efron and Tibshirani (1994). The stratum size, which is the number of first price auctions per sale, is 18 on average and ranges between 6 and 41.

Table 5 reports the scalar estimate overall and for two subgroups of bidders: those belonging to the top five and all others. Bidders in all groups substantially underbid. That is, holding rival bids constant, an individual bidder would optimize their expected payoff across auctions by increasing their own bid by a factor of more than three. Clearly, we can reject the null of BNE bidding in first-price auctions, suggesting that bidders coordinate to keep prices substantially below market value. Note that the test does not establish that tripling bids would result in BNE bidding as rival bids were held constant when computing α^*; if one bidder were to triple her bid, a profitable deviation may still exist for rivals.

Figure 2 plots the estimated per-auction profit function varying the bid scalar α. The left panel considers the top five, while the right panel considers all other bidders. The figures report 10th and 90th profit quantiles across bootstrap samples. In both cases profits are

Figure 2: Simulated Profit Varying the Bid Scalar α.

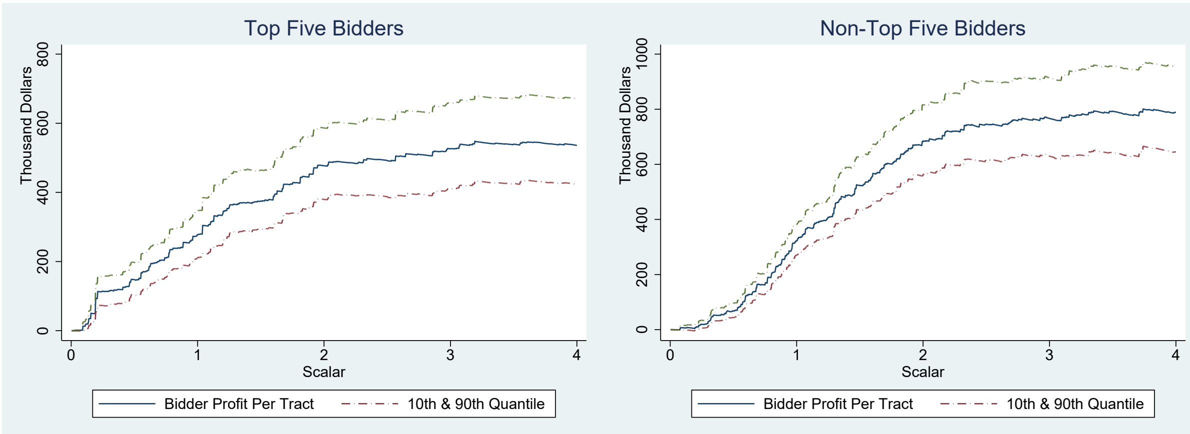

increasing at $\alpha = 1$, with the simulated optimal profit about twice as large as the observed profit.

Empirical auction work typically imposes assumptions on the information structure available to bidders and the details of the BNE. A key advantage of the underbidding test is that these details do not have to be specified. The test is robust to informational assumptions and the details of the BNE because it utilizes observable ex post oil extraction returns and bid data. Ex post returns are readily available in our setting and are available more broadly in settings where the lots can be resold or there is an active secondary market for the good.

While the evidence of underbidding rejects the hypothesis of BNE bidding, it does not address how far bonus bid payments are away from BNE bidding. The next section provides an answer to this question.

7 Minimum Revenue Bound

This section examines whether the observed revenues satisfy the robust lower revenue bound. The bound test is applicable regardless of the underlying information structure available to bidders and regardless of the BNE played. Any BNE of the first-price auction under any information structure must satisfy this bound.

The bound is obtained from the distribution of ex post returns and stems from the "worst-case" information structure and BNE that could arise. The intuition for the bound is based on the observation that the BNE bidding calculus in first-price auctions requires that bidders avoid losing the item to rival bidders.

An intuition for the bound can be developed by considering the asymmetric private values case. In the complete information case, the bidder with the high value wins the good and pays a price equal to the maximum of $N-1$ lowest values. The highest valuation of the losing bidders places a lower bound on the revenue. Next, suppose there is incomplete information such that bidders learn whether or not they have the high value and low value bidders do not learn which value they have. Low-value bidders expect their value to be the average of the $N-1$ lowest draws. The losing bidders will be willing to bid up to their expected value. Hence, the lower revenue is bounded based on the average of the $N-1$ lowest draws.

Bergemann et al. (2017) characterize the bound and show that it is obtained as a BNE when signals x_i are drawn from $G(X_i)^{1/N}$ and correlated with the value so that $\max x_i = v$. We do not postulate that the information structure in New Mexico oil sales takes this form but use it to simulate winning bids in the lower-bound BNE.

Characterization of the revenue bound requires the cdf of ex post returns. We estimate this cdf from our data on net returns v. Additionally, computation of the bound requires the number of potential bidders, which we estimate as the maximum across all auctions of the number of active bidders within an auction. While this is a commonly used approach in the empirical auction literature, we check the robustness to alternative definitions, including the minimally competitive case when the potential number of bidders equals 2. The third element is the reserve price, which we observe in the data. The reserve price is essentially linear in acreage until late in our sample period. For this reason, we normalize ex post returns, reserve prices, and bids by acreage. Thus, all our dollar value estimates are measure in per-acre terms.

The estimator proceeds in three stages: (i) estimate the potential number of bidders $\hat{N} = \max N^t$, (ii) estimate the cdf G using the empirical cdf $\hat{G}$, and (iii) estimate pseudo minimum revenues using the worst-case equilibrium.

The worst-case equilibrium has pseudo minimum revenues $\hat{\underline{\beta}}(v^t) = 0$ for values $v^t \leq \hat{v}$.[17] The threshold $\hat{v}$ is estimated as the largest ex post return v such that

$$\frac{1}{|V(\underline{v}, v)|} \sum_{x \in V(\underline{v}, v)} x \leq r^t,$$

where $V(u, w) = \{v^t : u < v^t \leq w\}$ denotes the set of ex post returns in the interval $(u, w]$. For $v^t > \hat{v}$ pseudo minimum revenues are estimated as:

$$\hat{\underline{\beta}}(v^t; r^t) \quad = \tfrac{1}{\hat{G}(v^t)^{(\hat{N}-1)/\hat{N}}} \left[r^t \cdot \hat{G}(\hat{v})^{(\hat{N}-1)/\hat{N}} + \tfrac{1}{|V(\hat{v}, v^t)|} \textstyle\sum_{x \in V(\hat{v}, v^t)} \left(\tfrac{\hat{N}-1}{\hat{N}} \cdot \tfrac{x}{\hat{G}(x)^{1/\hat{N}}} \right) \right] \quad .$$

Note that the estimator $\hat{\underline{\beta}}(v^t)$ is consistent when the data are iid. To see this, observe that the potential number of bidders $\hat{N}$ is a consistent and superefficient estimator. The empirical cdf is a consistent estimator of G by the Glivenko–Cantelli theorem. The sample average converges to the expected value by the law of large numbers. Hence, the threshold estimate $\hat{v}$ is consistent. Pseudo minimum bids $\hat{\underline{\beta}}(v^t)$ are consistently estimated as $\hat{G}$, $\hat{N}$ and $\hat{v}$ are consistent.

To test our null hypothesis

$$H_0 : \int_{v \in V} \left[\beta_w(x; r) - \underline{\beta}(x; r) \right] \, dG(x) \geq 0,$$

we use as our test statistic the sample average of the difference between observed auction

[17] Bergemann et al. (2017) show that minimum revenues $\underline{\beta}(v) = 0$ for $v \leq \hat{v}$ where the threshold $\hat{v}$ solves

$$\frac{\int_{\underline{v}}^{\hat{v}} x dG(x)}{G(\hat{v})} = r.$$

For values $v > \hat{v}$ the minimum revenues equal

$$\hat{\underline{\beta}}(v) = \frac{r \cdot G(\hat{v})^{(N-1)/N} + \int_{\hat{v}}^{v} \frac{N-1}{N} \cdot \frac{x}{G(x)^{1/N}} dG(x)}{G(v)^{(N-1)/N}}.$$

revenues and the inferred lower-bound auction revenues:

$$D_T = \frac{1}{T}\sum_{t=1}^{T}\left[\beta_w^t - \underline{\hat{\beta}}(v^t; r^t)\right].$$

We obtain the sampling distribution for the test statistic using a stratified bootstrap to account for correlation due to the exogenous variables z^t by resampling auctions inside strata (sale dates). Using a stratified bootstrap allows us to obtain a consistent measure of the sampling distribution for D_T by Assumption 1, see Efron and Tibshirani (1994). We use 199 bootstrap samples in the calculations.

Table 6 reports the test results. Our baseline specification, column "Pooled", assumes values v^t are iid, in which case bidders do not anticipate systematic variations in lease values over time. Our data indicate that this case is plausible as a regression of realized lease values v^t on a set of year fixed effects explains very little with an R-squared of 0.013. The null hypothesis that year fixed effects are jointly zero has a p-value of 0.024 and cannot be rejected at the 1 percent level. This case is also robust in the sense that it imposes minimal assumptions on the bidders' information structure, allowing them to range from fully informed to fully uninformed about tract values.

Columns "Annual" and "Monthly" in Table 6 report two alternative specifications in which we postulate a refinement of bidders' information structure by conditioning the distribution G on exogenous observables. Our data include on average 206 auctions per year and 18 auctions per month. We additionally estimate G_τ and $\underline{\hat{\beta}}_\tau$ for each year (or month) τ separately and report minimum revenue estimates by aggregating annual and monthly bound estimates. Since auctions within a month have an identical bid submission deadline, the monthly specification removes lease value uncertainty common to bidders. Note that if common uncertainty exists at the time of the auction, then the refined minimum revenue bound may be overstated.

Table 6 reports that on average across auctions the observed revenue equals 237 dollars

Table 6: Testing Revenues: $H_0 : D_T \geq 0$.

Period Length for G		Pooled	Annual	Monthly
Baseline Estimates	Realized Revenue	237.8	237.8	237.8
		(8.3)	(8.3)	(8.3)
	Minimum Revenue	783.6	878.1	1236.1
		(76.7)	(122.6)	(169.2)
	Test Statistic D_T	-545.9	-732.9	-998.3
		(76.5)	(122.0)	(169.6)
	P-Value	0.000	0.000	0.000
Robustness Checks				
(i) $\hat{N} = 2$	Minimum Revenue	409.1	459.3	648.9
		(40.1)	(64.0)	(88.4)
	Test Statistic D_T	-171.3	-221.5	-411.1
		(40.3)	(64.0)	(87.9)
	P-Value	0.000	0.000	0.000
(ii) Using Future Prices	Minimum Revenue	943.2	1105.1	1572.1
		(87.0)	(137.9)	(197.9)
	Test Statistic D_T	-705.4	-867.3	-1334.3
		(86.4)	(137.6)	(198.1)
	P-Value	0.000	0.000	0.000
(iii) Doubling Well Costs	Minimum Revenue	600.7	672.9	963.4
		(93.2)	(121.9)	(143.4)
	Test Statistic D_T	-362.9	-435.1	-725.6
		(93.3)	(122.2)	(143.3)
	P-Value	0.000	0.000	0.000
(iv) Values v Divided by 3	Minimum Revenue	262.6	297.6	415.9
		(26.1)	(41.76)	(56.5)
	Test Statistic D_T	-24.8	-59.8	-178.1
		(27.4)	(42.1)	(56.6)
	P-Value	0.1507	0.1507	0.000

Standard deviations are reported in parenthesis. The standard deviations of variables and the p-value of the null hypothesis are estimated using 199 bootstrap samples by resampling using the stratified bootstrap.

per acre measured in constant 2000 dollars. In contrast, the minimum revenue per auction ranges equals 783 dollars. The test statistic D_T is negative. We can reject the null of BNE bidding at all confidence levels. Revenues are less than one third of minimum revenues predicted by auction theory.

If value distributions are estimated separately every year and month, then the revenue bound increases to 878 and 1,236 dollars. The test statistic D_T remains negative in all cases. We can reject the null of BNE bidding at all confidence levels. If value distributions are estimated every month separately, then revenues are one fifth of predicted minimum revenues.

Robustness Checks

How robust is the test result? Table 6 reports robustness checks which relax one or more assumptions used to calculate the test statistic.

First, we consider reducing the potential number of bidders by replacing the consistent estimator $\hat{N} = \max N^t$ with $\hat{N} = 2$. This is the least competitive scenario in which there are only two firms competing for oil and gas leases. The minimum revenue bound falls but remains above the realized revenue. The test statistic D_T rejects the null of BNE bidding at all confidence levels. Even with only two bidders, the state of New Mexico should expect at least a 70 percent higher auction revenue.

Our second robustness check examined what happens if bidders did not form expectations about future prices correctly and instead evaluated future quantities with the currently expected future price, see Hendricks et al. (1987). Then the minimum revenue bound per auction increases to 943 dollars per acre in the baseline pooled test. Again, the null of BNE bidding is rejected at all confidence levels.

Third, we artificially double well cost estimates. Doing so results in minimum revenue bound per auction of at least 600 dollars per acre. Revenues remain less than one half of minimum revenues predicted by auction theory and the null of BNE bidding is rejected at all confidence levels.

Our fourth exercise examines what level of uniformly lower value estimates is required to rationalize the observed bids as a BNE. Table 6 shows that dividing all value realizations by 3 results in minimum revenue bound per auction of 262 dollars per acre in the pooled case. The test-statistic D_T rejects the null of BNE bidding at the 15 percent significance level. Lowering value estimates further results in a lowering of the test statistic D_T and we can no longer reject the null in the pooled and annual specification. We can conclude that dividing all values by three is required to rationalize the observed bids as a BNE in the pooled and annual specifications.

8 Conclusion

This paper documents evidence of systematic underbidding in oil and gas lease sales in New Mexico. Features of this market are favorable towards bidder collusion. Leases cover small homogeneous units of land and are awarded at regular time intervals at in-person auctions. The buyer's side is highly concentrated, with half of all leases won by only four bidders who know each other well and interact regularly.

Using the ex post value of leased tracts, we test for the presence of non-competitive bidding in three ways. First, we test whether bidder participation decisions are uncorrelated conditional on ex post returns and find statistically significant evidence of both positive and negative pairwise correlation. Second, we test whether bidders maximize their expected profit (holding rival strategies constant) and find that bidders could approximately double their expected profit by more than tripling each submitted bid. Finally, we test whether the average auction revenue are below the theoretical lower bound and find that the state's revenue from bonus bidding on oil and gas leases are less than one-third of the lower bound.

In the best-case BNE, which arises when bidders are fully informed (or not informed at all) about the ex post value of a lease, a standard Bertrand argument establishes that bids equal the expected value and extract all the rent. In this best-case scenario, bonus bidding achieves revenues of 1,753 dollars per acre - a seven-fold increase in revenues relative to the status quo.

We propose several steps NMSLO can take, some of which it has already taken, to combat low auction revenues and to move toward the best-case outcome. First, the NMSLO could raise the reserve price, which has occurred in recent years. Using information from prior production outcomes of neighboring tracts, the reserve price could be raised much further to a level close to the predicted lease value. Second, information about lease values, from geological studies and historic production data on neighboring tracts can be made available to bidders along with the lease sale announcement, which would reduce informational asymmetries between bidders and encourage competition. Third, NMSLO has made changes in regulation that make it more difficult for firms to acquire leases and renew them without drilling for oil. This makes the practice of hoarding land to protect any information rent a bidding ring may have more expensive, as it necessitates the drilling of wells. Fourth, barriers to entry could be reduced by attracting new bidders, which was encouraged with the shift to online auctions in 2016. Fifth, the identities of bidders could be concealed, making it more difficult to detect deviations from the collusive agreement. NMSLO introduced this practice when it moved to online auctions. Sixth, packages of tract can be offered at auctions instead of individual tracts and lease sales could take place at less frequent time intervals. Doing so will increase the benefits from deviation from a collusive agreement making collusion more difficult to sustain. Taken together, these steps may limit the potential for collusion which is a primary concern for the NMSLO if it is to meet its objective of "optimizing revenues while protecting [New Mexico's] heritage and [its] future."

References

Asker, J. (2010). A study of the internal organization of a bidding cartel. *American Economic Review 100*(3), 724–62.

Avery, C. (1998). Strategic jump bidding in english auctions. *The Review of Economic Studies 65*(2), 185–210.

Bajari, P. and L. Ye (2003). Deciding between competition and collusion. *The Review of Economics and Statistics 85*(4), 971–989.

Bergemann, D., B. Brooks, and S. Morris (2017). First–price auctions with general information structures: Implications for bidding and revenue. *Econometrica 85*(1), 107–143.

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.

Caoui, E. H. (2022). A study of umbrella damages from bid rigging. *The Journal of Law and Economics 65*(2), 239–277.

Chassang, S., K. Kawai, J. Nakabayashi, and J. Ortner (2022). Robust screens for noncompetitive bidding in procurement auctions. *Econometrica 90*(1), 315–346.

Chassang, S. and J. Ortner (2019). Collusion in auctions with constrained bids: Theory and evidence from public procurement. *Journal of Political Economy 127*(5), 2269–2300.

Conley, T. G. and F. Decarolis (2016). Detecting bidders groups in collusive auctions. *American Economic Journal: Microeconomics 8*(2), 1–38.

Efron, B. and R. J. Tibshirani (1994). *An introduction to the bootstrap.* Chapman and Hall, New York.

Green, E. J. and R. H. Porter (1984). Noncooperative collusion under imperfect price information. *Econometrica 52*(1), 87–100.

Harrington, J. E. (2005). Optimal cartel pricing in the presence of an antitrust authority. *International Economic Review 46*(1), 145–169.

Hendricks, K., J. Pinkse, and R. H. Porter (2003). Empirical implications of equilibrium bidding in first-price, symmetric, common value auctions. *The Review of Economic Studies 70*(1), 115–145.

Hendricks, K., R. H. Porter, and B. Boudreau (1987). Information, returns, and bidding behavior in ocs auctions: 1954-1969. *The Journal of Industrial Economics 35*(4), 517–542.

Hodgson, C. (2019). Information externalities, free riding, and optimal exploration in the uk oil industry. *Stanford SIEPR Working Paper No. 19-013 76*.

Igami, M. and T. Sugaya (2021). Measuring the Incentive to Collude: The Vitamin Cartels, 1990–99. *The Review of Economic Studies 89*(3), 1460–1494.

Kaplan, U., V. Marmer, and A. Shneyerov (2016). Identifying collusion in english auctions. *available at Semantic Scholar: Corpus ID: 53358781*.

Kawai, K. and J. Nakabayashi (2022). Detecting large-scale collusion in procurement auctions. *Journal of Political Economy 130*(5), 1364–1411.

Kawai, K., J. Nakabayashi, J. Ortner, and S. Chassang (2022). Using bid rotation and incumbency to detect collusion: A regression discontinuity approach. *The Review of Economic Studies, forthcoming*.

Kellogg, R. (2014). The effect of uncertainty on investment: Evidence from texas oil drilling. *The American Economic Review 104*(6), 1698–1734.

Kong, Y. (2020). Not knowing the competition: Evidence and implications for auction design. *RAND Journal of Economics 51*(3), 840–867.

Kong, Y. (2021). Sequential auctions with synergy and affiliation across auctions. *Journal of Political Economy 129*(1), 148–181.

Kong, Y., I. Perrigne, and Q. Vuong (2022). Multidimensional auctions of contracts: An empirical analysis. *American Economic Review 112*(5), 1703–36.

Krasnokutskaya, E. (2011). Identification and estimation of auction models with unobserved heterogeneity. *The Review of Economic Studies 78*(1), 293–327.

Laffont, J.-J. and Q. Vuong (1996). Structural analysis of auction data. *The American Economic Review 86*(2), 414–420.

Maskin, E. and J. Riley (1984). Optimal auctions with risk averse buyers. *Econometrica 52*(6), 1473–1518.

Matthews, S. A. (1983). Selling to risk averse buyers with unobservable tastes. *Journal of Economic Theory 30*(2), 370–400.

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

Miller, N. H., G. Sheu, and M. C. Weinberg (2021). Oligopolistic price leadership and mergers: The united states beer industry. *American Economic Review 111*(10), 3123–59.

Ordin, A. (2019). Investment and taxation: the case of oil and gas in the permian basin. *available at Semantic Scholar: Corpus ID: 210168158*.

Pesendorfer, M. (2000). A Study of Collusion in First-Price Auctions. *The Review of Economic Studies 67*(3), 381–411.

Porter, R. H. (1995). The role of information in u.s. offshore oil and gas lease auction. *Econometrica 63*(1), 1–27.

Porter, R. H. (2005). Detecting collusion. *Review of Industrial Organization 26*(2), 147–167.

Porter, R. H. and J. D. Zona (1993). Detection of bid rigging in procurement auctions. *Journal of Political Economy 101*(3), 518–538.

Porter, R. H. and J. D. Zona (1999). Ohio school milk markets: An analysis of bidding. *The RAND Journal of Economics 30*(2), 263–288.

Rotemberg, J. J. and G. Saloner (1986). A supergame-theoretic model of price wars during booms. *The American Economic Review 76*(3), 390–407.

Schurter, K. (2017). Identification and inference in first-price auctions with collusion. *available at Semantic Scholar: Corpus ID 31651513*.

Somaini, P. (2020). Identification in auction models with interdependent costs. *Journal of Political Economy 128*(10), 3820–3871.

Tamer, E. (2003). Incomplete simultaneous discrete response model with multiple equilibria. *The Review of Economic Studies 70*(1), 147–165.

Wilson, R. (1977). A bidding model of perfect competition. *Review of Economics Studies* (4), 511–518.