

Information Design in Common Value Auction with Moral Hazard: Application to OCS Leasing Auctions*

Anh Nguyen[†]

December 2024

Abstract

This paper explores the impact of information design on the auctioneer’s revenue in the US offshore oil/gas lease auctions where, post-auction, the winner decides whether to explore the auctioned tract and must pay the government a royalty on its production value. I first document that there is a positive correlation between the exploration rate and publicly observed losing bids. This suggests that the winning bidder uses the rivals’ bids to infer their private information about the tract’s potential. I then characterize the equilibrium bidding strategy when the auctioneer designs and commits to how to reveal information on losing bids to the winning bidder. Counterfactual exercises reveal that alternative bid disclosure policies significantly improve auctioneer revenue.

Keywords: Information Design; Common-Value Auction

JEL Classification: C57, D44, D82.

*I thank David Childers, Chris Conlon, Hanming Fang, Gagan Gosh, Karam Kang, Phillip Haile, Brent Hickman, Akshaya Jha, Robert Miller, Robert Porter, Paulo Somaini, Matthew Shum, Teck Yong Tan, Mo Xiao, and seminar and conference participants at Harvard University, Johns Hopkins University, Toulouse School of Economics, University of Michigan-Ann Arbor, University of Nebraska-Lincoln, University of Pennsylvania, University of Texas-Austin, Washington University in St. Louis, AEA Winter Meeting 2021, NBER SI 2021, IIOC 2021, and Stanford SITE 2022 for helpful comments and discussions.

[†]Tepper School of Business, Carnegie Mellon University. Email: anhnguyen@cmu.edu.

1 Introduction

In many common value auction settings, the winning bidder has to make investment decisions that depend on the expected value of the auctioned object. Often, this value remains uncertain to the winning bidder even after the auction. To address this uncertainty, the winning bidder utilizes the information contained in the bids of rival bidders to form a belief about the value of the auctioned object in addition to using its own private signal. In such cases, whether and how the losing bids are revealed impact not only the equilibrium bids but also subsequent strategic decisions and social welfare.

The often-studied Outer Continental Shelf (OCS) leasing auctions are an example of such a setting. In an OCS leasing auction, the auctioneer (the government) sells the right to explore and produce oil and gas on federal offshore tracts using a first-price sealed-bid format. After winning the auction, the winning bidder remains uncertain about the tract’s resource potential and must undertake costly exploratory drilling to resolve this uncertainty. The auctioneer’s revenue comprises both the winning bid and a royalty on the tract’s production revenue. If the tract remains unexplored, the royalty payment to the government is zero. Importantly, after the auction outcomes are determined, the losing bids are disclosed to the winning bidder. Given the common-value nature of the tract, lower losing bids can act as a discouraging signal for further exploratory drilling.

This paper studies how strategically withholding information about the losing bids can increase the auctioneer’s revenue in OCS auctions from 2000 to 2019.¹ Relative to a full disclosure benchmark in which all losing bids are disclosed to the winning bidder, the main economic tradeoff of an alternative disclosure policy is between lower bid revenue and potentially higher royalty revenue. The reduction in bid revenue is due to the increase in the uncertainty faced by the winning bidder when deciding whether to explore, resulting in a lower profit from winning the auction. However, by strategically pooling ‘bad’ information (i.e., low losing bids) with ‘good’ information (i.e., high losing bids), the auctioneer is able to increase the expected probability of drilling and hence royalty revenue.

¹Bid disclosure policies are utilized in practice in other auction settings, but there has been no systematic study of their effectiveness. For example, in 2009, the FDIC stopped releasing the identity of losing bidders and the second-highest bid in failed bank acquisition auctions, which sparked outrage from industry insiders (Fajt, 2009). In the US spectrum auction, the FCC utilizes a ‘limited disclosure’ rule in which only the number of bidders and the provisional winning bid are announced after each round (Xiao and Yuan, Forthcoming). Similarly, bidders on Google Ads auctions were concerned about the lack of transparency regarding information on losing bids (Joseph, 2020). In some cases, losing bidders prefer to withhold information on their bids. For instance, several US states blocked the release of information on their offers to host Amazon’s second headquarters.

To motivate this exercise, I first provide suggestive evidence that shows the significance of the aforementioned economic tradeoff. First, in this setting, the royalty revenue accounts for approximately 85% of the total revenue of the auctioneer. However, the exploratory drilling rate is low—only 24.5% of leased tracts were explored. This low exploration rate is an ongoing concern for the federal government, and multiple policy adjustments designed to improve the exploratory drilling rate have been proposed.² Second, the reduced-form analysis highlights a strong and positive correlation between the winning bidder’s exploration likelihood and the value of the losing bids. According to the regression results, a 1% increase in the average losing bid is associated with a 0.035 percentage-point increase in the probability of exploratory drilling. This suggests that the winning bidder utilizes the rival firms’ bids to infer their private information on the tract’s production value. As a comparison, a 1% increase in the winning bid is associated with a 0.039 percentage-point increase in the probability of exploratory drilling. As a firm’s bid is increasing in its expectation of the tract’s production value, this implies that the winning firm’s exploration decision is similarly affected by the (inferred) information of its rival firms as by its own private information.

I construct a model of a first-price sealed-bid common value auction where the auctioneer designs and commits to how the information on the losing bids is revealed. This disclosure policy is publicly announced to all bidders before bid submission. In the model, bidders have private information about the production value of a tract, which they use to form their bids, but they also face uncertainty regarding the cost of exploratory drilling. After the auction, the winning bidder receives information about the losing bids and the realized exploration cost, updates its belief about the tract’s profitability, and then decides whether to undertake exploratory drilling. My first contribution is to characterize the equilibrium bidding strategy in such a setting. Under a policy of full disclosure of the losing bids, the equilibrium bidding strategy resembles that of a standard common value auction, exhibiting the ‘*winner’s curse*’ phenomenon: winning the auction indicates that the winning firm is overly optimistic on the production value of the tract; anticipating this effect, every firm depresses its bid. Under a nondisclosure policy in which the winning bidder does not receive any information about losing bids, bidders further reduce their bids due to two effects. First, the expected value from winning the auction now decreases because of the lack of information available to the winning bidder *ex post*. Second, by not observing the second-highest bid, the winning bidder perceives the auction as being less competitive. I show that this intuition generalizes to the

²For example, the 1995 Deepwater Royalty Relief Act provided large royalty suspension volumes to lessees of deepwater tracts to promote exploration, development, and production.

case of a more general bid disclosure policy, and under a monotonic symmetric bidding equilibrium, the equilibrium bid is always lower than that under a full-disclosure policy.

While the theoretical prediction on the reduction of bid revenue when switching away from a full-disclosure policy is clear, whether this reduction can be offset by an increase in royalty revenue is an empirical question. My empirical analysis relies on two datasets. The first is a publicly available dataset on the US offshore auction outcomes, which are under a full-disclosure policy, and the tracts' drilling and production activities post-auction. The second dataset is a proprietary dataset on all the contracts between the oil/gas firms and the offshore rig company employed by the oil/gas firms to engage in drilling. The objects of interests are (1) the joint distribution of the firm's private signals and the true latent oil/gas potential of the tract and (2) the distribution of the exploration cost. In the absence of any unobserved heterogeneity in the tracts' values, the identification of the joint distribution of private signals and the latent value of the tract is straightforward when this latent value is observed ex post. In my setting, additional difficulty arises because the ex post production value of a tract is available only if the tract is explored, but the exploration decision is endogenous to the winning firm's belief after the auction and the unobserved exploration cost.

Therefore, another of my contributions is to provide an alternative way to identify the parameters of a common value auction without relying on the ex post value. Identifying the joint distribution of the firms' private signals and the latent tract's value is equivalent to identifying (i) the marginal distribution of the signals and (ii) the conditional mean of the latent value as a function of all bidders' realized signals. The former object is identified from the empirical bid distribution. The identification argument of the remaining objects—i.e., the conditional mean function and the distribution of the exploration cost—utilizes the variation in exploration decisions across bid realizations. Since this variation comprises the variation in both the conditional mean across bids and the ex post exploration costs, which are not observed, I disentangle these two objects using a cost shifter that affects the distribution of the exploration costs but does not influence the firm's posterior belief about a tract's resource potential. I show that when such a cost shifter exists, the objects of the model are nonparametrically identified. In the estimation, the cost shifter is constructed using rig costs, which are a significant component of the exploration cost, of past contracts associated with the nearby tracts. Thus, this cost shifter represents the rig market condition of a tract prior to its auction date.

My estimated parameters show that the OCS auctions are a favorable setting for imple-

menting an alternative bid disclosure policy to increase the auctioneer’s expected revenue. First, although bidders’ private signals are positively correlated, the correlation is low (0.11 and 0.12 for deep and shallow tracts, respectively). This implies that there is substantial uncertainty about the private information of rival firms, and the winning bidder’s posterior belief about the profitability of a tract is sensitive to the information on the losing bids. Second, I find that the relationship between the likelihood of exploratory drilling and the expected value of the tract under full disclosure is nonlinear. This nonlinearity arises due to both the option value of not drilling and the shape of the exploration cost distribution. As understood from the information design literature (Kamenica and Gentzkow, 2011), the nonlinearity suggests scope for revenue gains from strategically designing the bid disclosure policy.

In the first counterfactual exercise, I consider the effect of the nondisclosure policy. In this environment, the bidders decrease their bids by 1.13% for shallow tracts and by 4.23% for deep tracts. However, the post-auction exploration rate increases by 0.4 percentage points on average, which leads to a 1.61% and 2.37% increase in royalty revenue from shallow and deep tracts, respectively, resulting in a net increase in overall revenue of \$0.63 million and \$0.81 million for each type of tract. The differential effects between deep and shallow tracts reflect the differences in the characteristics of the tracts. For example, bidders have more precise information about shallow tracts (i.e., a more informed prior belief) than about deep tracts; therefore, the effect of a change in the bid disclosure policy on drilling outcomes as well as equilibrium bids tends to be smaller for shallow tracts. In addition, I show that the significant increase in royalty revenue results from increased exploration of more productive tracts. Furthermore, achieving a positive revenue gain from a nondisclosure policy requires the correlation between bidders’ signals to be sufficiently low. For instance, if the correlation exceeds 0.5, implementing a nondisclosure policy leads to a revenue loss on deep tracts for the auctioneer.

In the next counterfactual exercise, I explore the gains from using more general disclosure policies. A natural generalization that encompasses both nondisclosure and full-disclosure policies is a policy in which the auctioneer releases information on all losing bids below a preset threshold and withholds information on the exact magnitude of bids above this threshold. This policy is theoretically interesting because it represents the royalty-revenue-maximizing disclosure policy when drilling costs are binary (either high or low). I also consider a policy in which the k highest (e.g., the highest and the second-highest) losing bids are not revealed. This policy is of particular empirical interest because it has been

applied in other procurement settings.³ The results suggest that a nondisclosure policy is revenue-maximizing for most tracts within the class of disclosure policies considered.

Although the auctioneer’s revenue increases significantly when an alternative policy is implemented, it is important to examine whether such a policy induces welfare loss. In my model, the socially optimal outcome is achieved only when the royalty rate is zero and the winning bidder receives all available information before engaging in exploratory drilling. Given a positive royalty rate, which is a tax on the production revenue for the winning firm, the probability of exploratory drilling is always distorted downward under a full-disclosure policy. By removing some information on the losing bids ex post, we introduce another distortion to the firm’s drilling incentive. However, I show that deviating from a full-disclosure policy might result in overall welfare *gains*. For instance, at the current royalty rate, which averages 12% in the sample, the nondisclosure policy reduces the distortion caused by the positive royalty rate for approximately 80% of tracts, leading to significant welfare gains for deep tracts. This welfare gain arises because the positive impact of a nondisclosure policy on drilling is concentrated in more productive tracts. Therefore, a nondisclosure policy, on average, restores the drilling incentives for productive tracts but does not cause excessive drilling for unproductive tracts.

The remainder of this paper proceeds as follows. I first discuss the related literature in the next section. I then discuss the data and motivating evidence for information design on bids in Section 3. I present the model in Section 4 and provide the identification argument and estimation results in Section 5. The counterfactual exercises are in Section 6. Finally, I conclude in Section 7.

2 Literature Review

My paper adds to the empirical literature on auction design with ex post actions. Earlier works, such as Athey and Levin (2001), Lewis and Bajari (2011), and Bajari et al. (2014), study the equilibrium bidding strategy in scoring auctions in which each bid is an incentive contract. For example, Athey and Levin (2001) study timber auctions in which firms bid a per-unit price for each timber species, and bidders have private information about volumes of species. In these papers, however, there is no ex post uncertainty. More recently, Bhat-tacharya et al. (2022) studies auction design for onshore oil auctions in which they model the winning bidders observing the production value of the lease ex post and strategically

³See the discussion in Section 6.2.2.

delaying production in response to oil price uncertainty. In my OCS setting, the incentive to strategically delay production is not a first-order issue because the production duration of an offshore lease can be longer than 30 years. Instead, I focus on the effect of auction outcomes on the winning bidder’s belief about the tract’s latent production value and subsequent drilling decisions.

My paper also complements the literature on the role of information in altering an economic agent’s behavior in general and in auctions in particular. Since the work of Kamenica and Gentzkow (2011), there is a rapidly growing theoretical literature on how to optimally design information revelation policy to persuade economic agents; however, there is still limited attention on quantifying the potential benefit of information design in empirical markets. In the context of auctions, theoretically, Milgrom and Weber (1982) and Eső and Szentes (2007) study the auctioneer’s information disclosure policy when the auctioneer has access to exogenous signals that are affiliated with the bidders’ valuations. Empirically, Takahashi (2018); Allen et al. (2022), and Krasnokutskaya et al. (2020) study environments in which the bidders are uncertain about the scoring rule of an auction with multidimensional bids. In these papers, the uncertainty in the scoring rule stems from the auctioneer having private information on the project, and whether this information is revealed to the bidder has implications for the bidding strategies. In contrast, in this paper, I study the strategic decision of an auctioneer to disclose the rival bidders’ information that is privately revealed to the auctioneer through their sealed bids.

In addition, my paper contributes to the literature on the identification of auction models. Hendricks et al. (2003) and Athey and Haile (2007) establish the identification result for a standard pure common value when the ex post values are observed, utilizing the bid inversion method of Guerre et al. (2000). However, as previously mentioned, in my setting, the ex post values are not perfectly observed because only the explored tracts can be productive, and the exploration decision is endogenous to the latent production value. Similar to Somaini (2020), my identification argument requires a cost shifter that affects the distribution of the ex post drilling cost of the winning bidders; however, the role of the cost shifter is different here than in Somaini (2020). Whereas in Somaini (2020), bidders’ costs are privately observed and interdependent, and the cost shifters are bidder-specific, in my model, bidders do not have private information about their drilling costs, and the cost shifter is needed only to recover the unobserved distribution of drilling costs.

Finally, my paper contributes to the literature on the role of information in strategic exploration decisions within the oil and gas market. Porter (1995) and Hendricks and Porter

(1996) examine the determinants of exploratory drilling in U.S. offshore oil and gas auctions from 1954 to 1979, focusing on information spillovers among neighboring leases and the informational effects of auction outcomes. However, the market environment during this period differs significantly from mine, as most tracts (75%) were explored, and auctions only occurred if a tract was nominated by an oil or gas company. As a result, Hendricks and Porter (1996) did not find a positive correlation between the value of the losing bid and the exploration rate. Lin (2009, 2013), and Hodgson (2018) explore the strategic decision of oil and gas companies to delay exploration in order to free ride on information from neighboring leases' exploration outcomes. In contrast, my paper abstracts from the timing of exploration decisions and instead focuses on the impact of auction outcomes on whether exploration occurs and how this affects optimal auction design.

3 Data and Motivating Evidence

3.1 Institutional Setting and Data

I use the publicly available data from the Bureau of Ocean Energy Management (BOEM) on OCS auctions from 2000 to 2019.⁴ Each year, the BOEM holds sales in which the right to drill and extract the oil and gas from each available tract is auctioned. The BOEM data include the bid amounts of all bids submitted to each auction, the identity of the bidders, and whether the government accepts the highest bid as the winning bid. The data also include information on all well activities, such as when a well was drilled and when it was abandoned, and production activities of each well. If there is no drilling activity within a fixed lease term, which is between 5 and 10 years, the ownership of the tract reverts to the government. Otherwise, the lessee can continue exploration and production activities on the lease.

I supplement the BOEM data with the rig contract data from Rystad, which document all the rig contracts that were active at any point in the period from 2000 to 2019. Rigs are vital capital required for well drilling and are not owned by the oil and gas companies that are the tracts' leaseholders. Rigs are often contracted on a daily rate for a period between 3 months and a year. The data specify the tender date, the commencement date, the parties involved in the contracts, the daily rig rates, and the scope for each rig contract. Table 26 in the Appendix summarizes the average duration and costs of rig contracts.

⁴Hendricks and Porter (2014) provides a comprehensive overview of the OCS auctions and how they evolved over the period from 1954 to 2002.

During this period, all of the available unleased areas are available to be auctioned. The auctions' format is first-price sealed-bid with an announced minimum bid and a fixed royalty rate. The royalty rate determines the amount of post-production revenue that the winning bidder must pay to the government should production occur. The government might choose to reject the highest bid if it deems that there is insufficient competition for the lease. During the data period, the government rejects the highest bid only when there is one bidder (0.84% of all auctions were rejected).

Prior to the auction, potential bidders first acquire seismic information from a geophysical company, which is then processed by geophysicists at the oil and gas firms and will be further updated as new data arrive. The methods used by the bidders to process the data are proprietary. After the auction's outcome is revealed, the winning bidder decides whether to conduct exploratory drilling to determine if the tract contains profitable natural resources. If the outcome of an exploratory well is successful, the firm may decide to drill development wells to start production, and the result of this development stage is more certain than that of the exploration stage.

In this setting, firms are allowed to bid jointly, but large crude oil and gas producers face restrictions. These bidders cannot form joint bids, nor can they bid separately if they have agreements with other restricted bidders. Additionally, they are prohibited from making pre-bid agreements to transfer potential lease interests to individuals on the restricted joint bidder list. From 2000 to 2019, in auctions with at least two bidders, 25% of bids were joint, and 44% of auctions included at least one joint bid. Across all joint ventures, joint bids are more common for highly competitive and deeper tracts. However, among these, joint bids between large firms are more likely for shallower tracts, where uncertainty about oil and gas deposits is typically lower. Furthermore, joint winning bids among large firms are not associated with a higher ex post probability of discovering oil or gas. This suggests that firms primarily form joint bids to meet the capital requirements for bidding and, potentially, to reduce competition rather than to pool information among informed bidders.⁵ In Appendix F, I present a detailed analysis of firms' incentives to form joint bids.

Table 1 presents the summary statistics for auction outcomes and subsequent exploration and production activities from 2000 to 2019, focusing on auctions with at least two bidders. These auctions represent 23% of all auctions conducted during this period. There is substantial bid heterogeneity both across and within the auctions. For example, the average

⁵Hendricks and Porter (1992) examines joint bidding in OCS auctions from 1954 to 1979. The restriction on joint bidding was introduced in December 1975. During this period, capital constraints were also a likely explanation for joint bid formation.

Table 1: Summary statistics of 2000-2019 auctions with at least two bidders

Statistic	N	Mean	St. Dev.	Max	Min
First bid	2,510	5,089	11,733	157,111	10.10
Second Bid	2,510	2,067	6,180	84,391	4.24
Average losing bid	2,510	1,308	3,480	56,964	4.24
Number of bids	2,510	2.71	1.30	13	2
Fraction of explored tracts	2,510	0.25	0.43	1	0
Oil production	288	1,820	5,571	42,880	0
Gas production	288	11,498	27,151	253,421	0

Note: The bids are expressed in thousands of dollars. Oil production is measured in thousand barrels of crude oil, and gas production is measured in thousand MCF (1 MCF is equal to 1032 cubic feet). Some tracts produce only oil or gas, and some tracts produce both resources.

*In the empirical analysis, I utilize all available production data, which cover 1954 to 2019. Table 19 in the Appendix presents the summary statistics of this extended data period.

highest bid is more than twice as high as the average second bid. The exploration rate during this period is low: only 24.5% of leases were explored, among which 46.8% were subsequently developed.⁶ The low exploration rate in this period is mainly due to the change in the composition of available tracts which are now heavily concentrated in the deepwater area. Deeper tracts require higher costs of exploration and development, and their resource potential is inherently more uncertain. Figure 16 in the Appendix shows the increase in the average tract depth and the decline in the percentage of successfully developed tracts among explored tracts over time. In the Appendix, I also provide evidence that the low exploration rates are not due to oil/gas firms' binding exploration capacity constraints (Figure 17 and Table 18).

In the OCS setting, due to the spatial correlation in oil and gas deposits, the observed characteristics of a tract are high-dimensional. For example, the presence of a productive tract makes it more likely for surrounding tracts to also be productive. Therefore, the observed characteristics of a tract can include the relative distance of a tract to nearby explored tracts, the relative distance to nearby productive tracts, and the production quantities of those tracts. These characteristics affect both the equilibrium bidding strategies and the drilling probabilities. Figure 23 in the Appendix displays the spatial distribution of exploration outcomes (whether oil/gas was found) and the production quantities of productive tracts in my data.

⁶The average exploration rate of auctions with only one bidder is 8.9%.

To reduce the dimension of the observed characteristics of a tract, I construct a set of parameters to summarize the distribution of latent oil and gas quantities of each tract based only on the exploration and production outcomes of previously explored tracts. In this construction, oil and gas deposits are assumed to distribute in space according to a Gaussian random field. For each tract, the distribution of a deposit type (oil or gas) is assumed to be 0 with some random probability and follows a Gamma distribution otherwise. The scale parameters of the Gamma distributions are also random, whereas the shape parameter is a constant across tracts for each deposit type. The random components are assumed to be correlated between tracts with the degree of correlation being determined by the distance between tracts. The parameters of the Gaussian field are estimated using all available production data (Table 19 in the Appendix); however, the predicted oil/gas quantities of a tract are computed based only on the realized outcomes of previously explored tracts. Further details on this construction appear in Appendix E, and the estimated parameters of the Gaussian random fields and the shape parameters of the Gamma distributions are reported in Table 40 in the Appendix .

Table 2: Distribution of $(\mu_{Bernoulli,oil}, \mu_{Bernoulli,gas}, \mu_{Gamma,oil}, \mu_{Gamma,gas})$ across tracts

Statistic	N	Mean	St. Dev.	Max	Min
Bernoulli Prior Mean (Oil)	2,510	0.11	0.11	0.53	0
Bernoulli Prior Mean (Gas)	2,510	0.25	0.28	0.90	0
Gamma Prior Mean (Oil)	2,510	22,375	22,787	107,155	22.87
Gamma Prior Mean (Gas)	2,510	53,361	25,042	358,667	7,167

Note: the mean column is the average of the estimated parameters across all tracts in the sample of tracts that were auctioned between 2000 and 2019. The standard deviation is also across tracts. The standard errors of these estimates are not reported. $\mu_{Gamma,oil}$ is measured in thousand of barrels of oil, and $\mu_{Gamma,gas}$ is measured in thousand of MCF.

Under this construction, each tract is now characterized by six parameters $(\mu_{Bernoulli,oil}, \mu_{Bernoulli,gas}, \mu_{Gamma,oil}, \mu_{Gamma,gas}, var_{Gamma,gas}, var_{Gamma,oil})$ where $\mu_{Bernoulli,oil}(\mu_{Bernoulli,gas})$ is the probability that the tract contains oil (gas), $\mu_{Gamma,oil}(\mu_{Gamma,gas})$ is the predicted quantity of oil (gas) if oil (gas) is found, and $var_{Gamma,oil} (var_{Gamma,gas})$ is the variance of the distribution of oil (gas) if oil (gas) is found. Because the shape parameter of the Gamma distributions is the same for all tracts, the variance of the Gamma distributions is proportional to the squared value of the means scaled by the estimated shape parameter. Therefore, in the subsequent analysis, I restrict the set of tract's predicted production vari-

ables to $(\mu_{Bernoulli,oil}, \mu_{Bernoulli,gas}, \mu_{Gamma,oil}, \mu_{Gamma,gas})$. These constructed values across tracts are summarized in Table 2. The table shows substantial variation across tracts and a low ex ante probability of finding any oil or gas. For example, based on these estimates, a tract in the sample, on average, has a 11% probability of containing oil and a 25% probability of containing gas. The constructed predicted production variables explain 82.95% (34.28%) of the variation in whether gas (oil) is found, and 17.08% (41.08%) of the variation in gas (oil) production quantity for productive tracts.

Using a similar procedure as described above, but incorporating the realized outcomes of all tracts (rather than only previously explored tracts), I also construct a measure of the ex post belief on oil and gas potential. The ex post belief reflects all information about a tract based on the outcomes of all explored tracts prior to 2019. For productive tracts, the ex post belief equals the realized production value. For tracts that have been explored but found unproductive, the ex post belief equals zero oil and gas discovered. For unexplored tracts, the ex post belief is more informative than the predicted production variables because it incorporates more precise data from neighboring tracts. Figure 24 in the Appendix illustrates the differences between the predicted production variables and the constructed means of the ex post belief.

3.2 Motivating Evidence

In this paper, I focus on the extent to which the negative informational content from the auction outcomes influences the drilling decision of the winning bidder. Columns (1)–(4) of Table 3 show a positive and significant correlation between exploration decisions and the value of the average losing bids, and this correlation persists even after including various control variables. For instance, in columns (2) and (4), I incorporate variables characterizing the predicted oil and gas production of each tract $(\mu_{\Gamma,oil}, \mu_{\Gamma,gas}, \mu_{Bernoulli, oil}, \mu_{Bernoulli, gas})$.⁷ In columns (3) and (4), I include controls for neighboring tracts’ characteristics, such as whether a neighboring tract was explored and whether this occurred before or after the auction.⁸ I also control for the winning bidder’s characteristics, including whether the winning bidder owns neighboring leases and whether these leases have been explored and found productive.

⁷The results in column (2) show that these predicted production variables are strong predictors of exploration decisions. For instance, a 1 percentage point increase in the predicted probability of finding oil corresponds to a 0.3 percentage point increase in the likelihood of exploration. Moreover, compared to the results in column (4), the findings from column (2) suggest that these variables capture significant heterogeneity across tracts in terms of production potential, exploration likelihood, and location characteristics.

⁸For each tract, I define “neighboring tracts” using a metric similar to that in Hendricks et al. (2003), identifying tracts within 0.11 degrees of latitude and 0.12 degrees of longitude.

Furthermore, the magnitude of the effect of the average losing bid on the exploration rate is similar to that of the winning firm’s own bid, suggesting that the winning firm’s exploration decision is similarly affected by the inferred information of its rivals as by its own private information.⁹

The regression results above may suffer from omitted variable bias if bidders have access to information about a tract’s oil and gas potential beyond what is reflected in their bids. For example, the winning bidder might conduct additional surveys and base their exploration decisions on newly acquired information that was not part of the initial bid formation. However, the data suggest that ex post information acquisition is rare. Only 10% of tract owners applied for additional permits for 3D and 4D surveys—considered the most informative for exploration and development—before a tract was explored (Table 21 in the Appendix). This 10% statistic likely represents an upper bound, as it includes permit applications for other not-yet-auctioned tracts by the same lessees. To further address this concern, in Column (5) of Table 3, I include variables that characterize ex post beliefs in the regression. This approach mitigates omitted variable bias concerns under the assumption that unobserved post-auction information correlates with the ex post variables and that the measurement error from using these ex post variables instead of the true latent information is uncorrelated with the losing bids. The results indicate that the positive correlation between losing bids and exploration likelihood persists even after accounting for ex post information.¹⁰

In my analysis, I assume that each winning bidder considers its exploration decision independently between tracts. However, because the exploration outcomes on the adjacent tracts are usually positively correlated, the leaseholders of the adjacent tracts might choose to explore cooperatively to mitigate potential losses from overdrilling inefficiencies (Hendricks and Porter, 2014). To measure the extent to which firms cooperate on their exploration decisions among neighboring leases, I exclude explored (unexplored) tracts whose owners entered into ownership agreements with nearby unexplored (explored) tracts, which account

⁹In Table 35 in the Appendix, I provide an alternative regression specification that separately examines the effects of the second-highest bid, third-highest bid, and fourth-highest bid on the exploration decision. The regression results indicate that higher losing bids are associated with a greater likelihood of exploration. However, the effect of the fourth-highest bid is imprecise due to the limited number of auctions with four or more bids.

¹⁰The increase in R^2 when including the ex post variables arises from the construction of the ex post Gamma variables, which are equal to zero only if a tract is explored and found unproductive; otherwise, these values are strictly positive for productive or unexplored tracts. Excluding the ex post Gamma variables results in R^2 values that are similar to those obtained when the ex post variables are not included, indicating that ex post outcomes—such as whether oil or gas is found—do not explain the variation in exploration rates (Table 32 in the Appendix).

Table 3: Correlation between the exploratory drilling rate and the auction outcomes

	Exploration drilling = TRUE					
	(1)	(2)	(3)	(4)	(5)	(6)
Winning bid (in log)	0.059 (0.010)	0.055 (0.010)	0.048 (0.010)	0.048 (0.010)	0.039 (0.008)	0.036 (0.009)
Average losing bid (in log)	0.040 (0.014)	0.051 (0.014)	0.041 (0.013)	0.042 (0.013)	0.035 (0.012)	0.026 (0.013)
Number of bids	0.018 (0.008)	0.015 (0.008)	0.013 (0.008)	0.012 (0.008)	0.007 (0.007)	0.003 (0.008)
Predicted production for oil (Bernoulli)		0.311 (0.112)		-0.048 (0.134)	-0.260 (0.391)	-0.016 (0.460)
Predicted production for gas (Bernoulli)		0.405 (0.062)		0.258 (0.122)	0.422 (0.259)	0.241 (0.287)
Predicted production for oil (Gamma) (in log)		-0.014 (0.011)		0.032 (0.018)	0.053 (0.016)	0.039 (0.016)
Predicted production for gas (Gamma) (in log)		-0.024 (0.022)		-0.022 (0.025)	0.008 (0.022)	0.026 (0.022)
Ex post mean for oil (Bernoulli)					0.501 (0.361)	0.277 (0.427)
Ex post mean for gas (Bernoulli)					-0.362 (0.283)	-0.168 (0.312)
Ex post mean for oil (Gamma) (in log)					-0.039 (0.002)	-0.040 (0.002)
Ex post mean for gas (Gamma) (in log)					-0.036 (0.002)	-0.039 (0.003)
Neighbor and winning bidder characteristics	No	No	Yes	Yes	Yes	Yes
Area FE	No	No	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Exclude Shared Ownership	No	No	No	No	No	Yes
Observations	2,510	2,510	2,510	2,510	2,510	1,988
Adjusted R ²	0.197	0.240	0.315	0.316	0.473	0.539

Note: The variables that describe the characteristics of winning bidders include the number of tracts owned by lessees that have not been explored or under production, those under explored, and those under production. The variables that describe the characteristics of neighboring tracts include the average bids from neighboring tracts, the number of neighboring tracts that were explored, and the number of neighboring tracts that were under production. In Column (6), leases with shared ownership of nearby tracts are excluded. Standard errors are in parentheses and adjusted for heteroskedasticity. The ex post Gamma variables may equal zero, so they are scaled up by adding 1 to ensure that the log value is well-defined.

for 13.2% of leases in the sample. The qualitative results of the regressions remain the same (Columns (6) of Table 3).¹¹

¹¹When the exploration decision is not independent between tracts, the bidding behavior will also be interdependent. In my model, I account only for the effect of the realized production outcomes of neighboring tracts on bidding behavior. Kong (2021) studies a model in which both the affiliation between tracts and the potential synergies (for example, due to economies of scale in drilling activities) in a private value paradigm

In the subsequent analysis, I assume that bidders are symmetric in their ability to process seismic information. Consequently, the information contained in each bid is equally informative to the winning bidder. In Figure 20 in the Appendix, I show that, although there are across-firm differences in success rates among tracts that were explored, these differences do not seem to stem from the differences in firm’s size or experience. Instead, they are likely due to sampling variation. For instance, consider Chevron (formerly known as California Oil Company) and Shell Offshore, Inc—both significant market players in terms of the number of tracts explored and the total bids placed. Chevron successfully explored approximately 25% of its tracts, while Shell’s success rate was about 40%, both figures falling below the average success rate of 50% as presented in Table 1. This conclusion holds even when the auction’s and the tract’s characteristics are considered.

4 Model

I first present a model of a common value auction with ex post moral hazard under full revelation of the bids ex post. This benchmark model is also the model used for identification and estimation. I then discuss an extension of the model in which the auctioneer does not disclose the losing bids. Throughout this section, I abstract from some empirically relevant details, which will be introduced in Section 5.

4.1 Full-Disclosure (FD) Model

The federal government (auctioneer) wants to auction off the lease of an offshore tract via a first-price sealed-bid auction. There are $N \geq 2$ firms/bidders. N is common knowledge. The expected production value during the lease term is a random variable $Q \in \mathbb{R}_+$, and the common prior belief about Q is represented by the distribution F^Q with density f^Q . Prior to the auction, each firm i receives a *private* signal $S_i \in \mathbb{R}_+$ about Q , which represents each firm’s private information about the potential production value of the tract. Let $\mathbf{S} = (S_1, \dots, S_N)$ and $f^{\mathbf{S}}(\mathbf{S}|q)$ be the joint density of the firms’ signals conditional on $Q = q$. I assume that for all $q \in \mathbb{R}_+$, $f^{\mathbf{S}}(\mathbf{S}|q)$ has full support.¹² Let the expected production value

can arise.

¹²Since $f^{\mathbf{S}}(\cdot|q)$ has full support, by Bayes rule, $f^Q(q|\mathbf{S} = \mathbf{s}) = \frac{f^{\mathbf{S}}(\mathbf{S}|q)}{\int f^{\mathbf{S}}(\mathbf{s}|q')f^Q(q')dq'}$.

conditional on $\mathbf{S} = \mathbf{s}$ be denoted by

$$V(\mathbf{s}) := \int q \frac{f^{\mathbf{S}}(\mathbf{s}|q)f^Q(q)}{\int f^{\mathbf{S}}(\mathbf{s}|q')f^Q(q')dq'} dq$$

where I assume that $V(\cdot)$ is differentiable almost everywhere and is symmetric in its arguments. Without loss of generality, I also assume that $V(\mathbf{s})$ is increasing in s_1 to s_N and normalize the signals such that for $\mathbf{s} = (0, \dots, 0)$, $V(\mathbf{s})$ is 0. I make the standard assumption that S_1, S_2, \dots, S_N are affiliated.¹³

Besides the winning bid, the government also charges a royalty rate of $r < 1$ on the tract's production value, where r is announced before the auction. To extract the natural resources from the tract, the firm must incur an exploration cost. This cost is realized only *after* the auction—for example, such exploration costs depend on the geophysical and geological characteristics of the tract and can be fully assessed by a firm only after it has won the auction and gained access to the tract.¹⁴ We let $C > 0$ denote the random variable of the exploration cost and F^C be its distribution. F^C is also common knowledge.

To summarize, suppose that firm i wins the auction with a bid of b . Under the realizations $Q = q$ and $C = c$, if firm i explores the tract, then its overall profit is $(1 - r)q - c - b$, and the government's profit is $b + rq$. On the other hand, if firm i does not explore, its profit is $-b$, whereas the government's profit is b . The game proceeds as follows:

1. Each firm i *privately* observes the realization of S_i .
2. All the firms simultaneously submit a sealed bid to the auctioneer.
3. The government announces the winner, who is the highest bidder; the winner then pays its bid. In the event of a tie, the winner is randomly chosen from among the highest bidders.
4. The auctioneer reveals all the submitted bids.
5. The winner observes the realization of the exploration cost C .

¹³That is, for any $\mathbf{s} = (s_1, \dots, s_N)$ and $\mathbf{s}' = (s'_1, \dots, s'_N)$,

$$f^{\mathbf{S}}(\mathbf{s} \vee \mathbf{s}') f^{\mathbf{S}}(\mathbf{s} \wedge \mathbf{s}') \geq f^{\mathbf{S}}(\mathbf{s}) f^{\mathbf{S}}(\mathbf{s}'),$$

where $\mathbf{s} \vee \mathbf{s}'$ and $\mathbf{s} \wedge \mathbf{s}'$ are the component-wise maximum (i.e., join) and minimum (i.e., meet) between \mathbf{s} and \mathbf{s}' , respectively.

¹⁴If the exploration cost is known with certainty before the auction, in equilibrium, all bidders with a positive bid will explore the tract after the auction. This is because a firm will only participate in the auction of the expected valuation of a tract is higher

6. The winner decides whether to explore the tract. If the winner decides to explore, it incurs C , and the value of Q is realized. Subsequently, the firm collects Q and pays rQ to the government. On the other hand, if the firm decides not to explore, the game ends.

In this model, I assume that bidders are uncertain about exploration costs before bidding. This is because exploration costs are primarily driven by rig market conditions, rig availability, and the technical challenges associated with drilling, which determine the duration of the operation. The uncertainty in exploration costs arises from two main sources. First, potential seafloor geologic and manmade hazards related to drilling activities can typically only be assessed closer to the actual drilling date. Specifically, a required site-specific high-resolution geohazard survey at the exact location of the planned exploration well is usually conducted just the season before a drilling rig is mobilized to the site (BOEM, 2019). Second, there is substantial variation in rig contract rates, even among rigs contracted by the same oil and gas company, as these rates are determined by market conditions where both sides of the market are unconcentrated (Kaiser et al., 2013; Vreugdenhil, 2020). Figure 29 in the Appendix illustrates the variation in contract rates for each oil and gas company in my data.

Since the uncertainty in exploration costs is unlikely to be related to the underlying production value of a tract, I further assume that the realization of exploration costs in Stage 5 does not alter the winning bidder’s posterior belief about the tract’s production value. This is stated in Assumption 1.¹⁵

Assumption 1 C and Q are independent.

Another implicit assumption of the model is that bidders do not possess private information about exploration costs. The fact that tracts that were previously unexplored and whose leases have expired often receive no bids upon re-auctioning suggests that cost heterogeneity among bidders is not the primary factor driving the low rate of exploration.¹⁶

¹⁵While exploration costs are unlikely to be correlated with a tract’s production value, development costs are likely to be, as the size of the deposit determines the number of development wells that need to be constructed. In Section 5, I assume that development costs are a multiplicative function of the production value.

¹⁶This assumption is also made for tractability. In a model with both common value and private value components (possibly correlated), the problem becomes an auction in which each bidder holds multiple signals. Without imposing strong parametric assumptions on how these signals combine into a single sufficient statistic, it is difficult to derive the equilibrium bidding strategy.

4.1.1 Equilibrium Analysis

Each firm's strategy consists of a bidding function $\beta_i : S_i \rightarrow \mathbb{R}_+$ and an exploration decision (conditional on winning the auction) that depends on the realization of C and all the information about the expected production value after the auction. The winner of the auction will explore if and only if its expected value of $(1-r)Q$ conditional on all the available information is higher than the realization of C . Given the simplicity of the drilling decision, it suffices to focus on only the equilibrium bidding functions in the equilibrium analysis. Henceforth, I abuse terminology and refer to an equilibrium as only the profile of the bidding functions of the equilibrium strategies. I restrict attention to symmetric Bayes Nash equilibria in which each firm i bids according to the same bidding function $\beta : S_i \rightarrow \mathbb{R}_+$, where $\beta(\cdot)$ is strictly increasing.

Let Y_i denote the first-order statistics of the $N-1$ random variables S_j for $j \neq i$ (i.e., $Y_i = \max_{j \neq i} S_j$). Let $G^Y(Y_i|s)$ denote the distribution of Y_i conditional on $S_i = s$, and let $g^Y(Y_i|s)$ be the corresponding density. Let \mathbf{Z}_i denote the $(N-2)$ -dimensional random vector $(Z_i^{(2)}, Z_i^{(3)}, \dots, Z_i^{(N-1)})$, where $Z_i^{(k)}$ denote the k -order statistics. Let $G^{\mathbf{Z}}(\mathbf{Z}_i|s, y)$ denote the distribution of \mathbf{Z}_i conditional on $S_i = s$ and $Y_i = y$. Let $g^{\mathbf{Z}}(\mathbf{Z}_i|s, y)$ be the corresponding density. Finally, let $G^{Y, \mathbf{Z}}(Y_i, \mathbf{Z}_i|s)$ denote the joint distribution of (Y_i, \mathbf{Z}_i) conditional on $S_i = s$, and let $g^{Y, \mathbf{Z}}(Y_i, \mathbf{Z}_i|s)$ be the corresponding density.

Next, let $\psi(V)$ denote the net revenue from the tract when its ex post expected production value is V , taking into account the firm's drilling decision.

$$\psi(V) := \int_0^\infty \max\{0, (1-r)V - c\} dF^C(c) = \int_0^{(1-r)V} ((1-r)V - c) dF^C(c) \quad (1)$$

I make two observations. First, $\psi(V)$ is nonlinear in V because of the option value of not drilling. Because of this nonlinearity, changes in the observability of V affect the bidder's expected net revenue and hence actions, both at the drilling and the bidding stages, making bid disclosure policy a useful market design tool. Second, the curvature of $\psi(V)$ depends on the cost distribution F^C . To see this, note that $\psi'(V) = (1-r)F^C((1-r)V)$.

Let $v(s, y)$ denote the expected net revenue conditional on $S = s$ and $Y = y$ when the winning bidder observes the realization of \mathbf{S} before drilling.

$$v(s, y) := \int_{\mathbf{z}} \psi(V(s, y, \mathbf{z})) dG^{\mathbf{Z}}(\mathbf{z}|s, y) \quad (2)$$

Therefore, $v(s, y)$ is also firm i 's expected profit from winning the auction conditional on

$S_i = s$ and $Y_i = y$, where the expectation is taken at the start of Stage 2. In deriving $v(\cdot)$, the assumption that all firms are playing according to an increasing strategy $\beta(\cdot)$ is utilized; therefore, observing all losing bids is the same as observing all private signals of losing firms.

Consequently, at the start of Stage 2, if $S_i = s$ and firm i conjectures that all the other firms bid according to $\beta(\cdot)$, firm i 's expected overall profit from bidding b is

$$\int_0^{\beta^{-1}(b)} [v(s, y) - b] g^Y(y|s) dy \quad (3)$$

Let $L(s'|s) := \exp\left(-\int_{s'}^s \frac{g^Y(t|t)}{G^Y(t|t)} dt\right)$.

Proposition 1 *There exists a unique increasing and symmetric Bayes Nash equilibrium. The equilibrium strategy, denoted by $\beta^{FD}(\cdot)$, is as follows*

$$\beta^{FD}(s) = \int_0^s v(s', s') dL(s'|s). \quad (4)$$

The proof of Proposition 1 is found in Appendix A. The proof here is similar to the standard common value auction (Milgrom and Weber, 1982). In the proof, I also show that, for any s , $L(s'|s)$ is a proper distribution that is first-order stochastically dominated by the distribution $G^Y(s'|s' \leq s)$. Therefore, $\beta^{FD}(s) < \mathbb{E}[v(s', s')|s' \leq s]$ for all $s > 0$. This bidding strategy thus reflects the firm's understanding of the "winner's curse" — winning the auction indicates that the firm is overly optimistic about the tract's production value; anticipating this effect, the firm depresses its bid.

4.2 Nondisclosure Model

I now consider the equilibrium bidding and drilling strategies when the auctioneer announces at the start of the auction that the losing bids will no longer be revealed, i.e., a nondisclosure (ND) policy. The modification to the baseline game is that Stage 4 is now removed, and this is common knowledge.¹⁷

4.2.1 Drilling Decision: FD vs. ND

Before turning to the equilibrium analysis of the bidding strategy, I first discuss the difference in the drilling decision under FD and ND. Conditional on bidder i 's signal being $S_i = s$,

¹⁷In addition to the ND policy, my model can be extended to incorporate more general disclosure policies. The equilibrium analysis of such policies is provided in Section 1 of the Online Appendix.

consider the winner's expected drilling probability at the end of Stage 3—i.e., before the realization of the cost of drilling.

Under an FD policy, the winning bidder's expected drilling probability after winning the auction but before the realization of its cost of drilling is given by

$$\Delta^{FD}(s) = \mathbb{E} [F^C(V(\mathbf{S})) | S_i = s, S_{-i} \leq s] \quad (5)$$

Under an ND policy, winning bidder i does not observe S_{-i} . Therefore, the expected drilling probability is given by

$$\Delta^{ND}(s) = F^C(\mathbb{E}[V(\mathbf{S}) | S_i = s, S_{-i} \leq s]) \quad (6)$$

Although the bidders are assumed to be risk neutral, when $F^C(\cdot)$ is nonlinear, $\Delta^{FD}(s)$ and $\Delta^{ND}(s)$ are generally different. The difference between these two objects can be readily understood by treating $F^C(c)$ as a utility function. Under ND, conditional on S , the bidder receives deterministically the value $\mathbb{E}[V(\mathbf{S}) | S_i = s, S_{-i} \leq s]$. Under FD, conditional on S , the bidder receives a lottery with the payoff distribution being determined by the distribution of $S_{-i} | S_i = s$. Therefore, when $F^C(\cdot)$ is always concave (convex), the drilling rates under an ND policy are always greater (smaller) than those under an FD policy, using Jensen's inequality. When $F^C(\cdot)$ admits both concave and convex regions, the theoretical predictions on the relative magnitude of Δ^{ND} and Δ^{FD} are indeterminate. The effect of an ND policy on the expected drilling probability depends on two factors: (1) the shape of the cost distribution $F^C(\cdot)$, and (2) the distribution of the posterior mean $V(\mathbf{S} | S_i = s, S_{-i} \leq s)$ for different realizations s .

Figure 4: Illustration of the effect of an ND policy on the drilling probability




Figure 4 illustrates this intuition when there are two bidders. I consider three scenarios. In the first scenario, the winning signal is s_1 , the support of the posterior mean is $[V(\mathbf{S}|S_i = s_1, S_{-i} = 0), V(\mathbf{S}|S_i = s_1, S_{-i} = s_1)]$, and $F^C(\cdot)$ is concave on this interval. Therefore, $\Delta^{ND}(s_1) > \Delta^{FD}(s_1)$. Conversely, in the second scenario, the winning signal is s_2 , and $F^C(\cdot)$ is convex on the interval $[V(\mathbf{S}|S_i = s_2, S_{-i} = 0), V(\mathbf{S}|S_i = s_2, S_{-i} = s_2)]$. As a result, $\Delta^{FD}(s_2) < \Delta^{ND}(s_2)$. In the last scenario, the winning signal is s_3 , and $F^C(\cdot)$ is neither concave nor convex over the support of the posterior mean. Therefore, the effect of ND relative to that of FD at the winning signal s_3 depends the distribution of the posterior mean $V(\cdot)$ over its support.

4.2.2 Equilibrium Analysis

In the following analysis, I explain the difference in the equilibrium bidding strategy between the FD and ND auctions.

As before, I focus on symmetric Bayes Nash equilibria in which each firm i bids according to a strictly increasing function $\beta^{ND} : S_i \rightarrow \mathbb{R}_+$. Let

$$\hat{V}_0(s, y) := E[V(S_i, Y_i, \mathbf{Z}_i) | S_i = s, Y_i \leq y] \quad (7)$$

$$\bar{V}_0(s, y) := E[V(S_i, Y_i, \mathbf{Z}_i) | S_i = s, Y_i = y] \quad (8)$$

Therefore, $\hat{V}_0(s, y)$ represents firm i 's expectation of Q conditional on $S_i = s$, and $S_{-i} \leq y$, and $\bar{V}_0(s, y)$ represents firm i 's expectation conditional on more granular information that $\max_j S_j = y$. The affiliation property of the private signals implies that $\hat{V}_0(s, y) \leq \bar{V}_0(s, y)$. I adopt the same notations as in Section 4.1


Proposition 2 *In the ND auction, the increasing and symmetric Bayes Nash equilibrium is uniquely*

$$\beta^{ND}(s) = \int_0^s \left(\psi(\hat{V}_0(s', s')) + \left[\psi'(\hat{V}_0(s', s')) (\bar{V}_0(s', s') - \hat{V}_0(s', s')) \right] \right) dL(s'|s). \quad (9)$$

Moreover, $\beta^{ND}(s) < \beta^{FD}(s)$ for all s —i.e., the equilibrium bids in the ND auction are lower than those in the FD auction.

The proof of Proposition 2 is in Appendix A. The difference between $\beta^{ND}(\cdot)$ and $\beta^{FD}(\cdot)$ in Eq. (4) is illustrated in Figure 5. First, recall that $v(s', s') = \mathbb{E}[\psi(V(\mathbf{S})) | S_i = s', Y_i = s']$ is firm i 's expected profit from winning the FD auction when conditioned on $S_i = Y_i = s'$.

Figure 5: Illustration of the difference between $\beta^{FD}(\cdot)$ and $\beta^{ND}(\cdot)$



In the ND auction, firm i 's expected profit from winning when conditioned on $S_i = Y_i = s'$ is $\psi(\bar{V}_0(s', s'))$. Note that $\psi(\cdot)$ is convex; therefore, $\psi(\bar{V}_0(s', s')) \leq v(s', s')$. The difference between these two terms (the red difference in Figure 5) reflects the anticipated decrease in revenue resulting from less information being available to the winning bidder in the ND auction.

The difference between $\psi(\bar{V}_0(s', s'))$ and the term inside the integral of $\beta^{ND}(\cdot)$ in (9) represents the effect of the unobservability of the second-highest bid in the ND auction. This is represented by the blue difference in Figure 5. Intuitively, the winning bidder now places more weight on the possibility that its closest rival will place a lower bid; therefore, this effect resembles the effect of lessening competition in the auction.

In summary, the ND auction aggravates the winner's curse relative to the FD auction, and the equilibrium bids are always lower in the ND auction. Intuitively, the value of the object being auctioned is determined by the latent production value of the tract and the available information about this latent value. Since information can always be ignored, bidders are always weakly worse off ex post when such information is unavailable. Therefore, the auctioneer faces a tradeoff when concealing the losing bids. On the one hand, the ND auction yields strictly lower bid revenue than the FD auction. On the other hand, the expected drilling rate may be higher in the ND auction. Thus, it is unclear which of the two auctions yields the auctioneer a higher expected total profit.¹⁸

¹⁸As a referee pointed out, the use of an ND policy may also affect bidders' incentives to form joint bids. In Section F.1 of the Appendix, I use a simple parametric example to explore this and demonstrate that bidders are *less* likely to form joint bids under an ND policy. This result follows the intuition that if bidders do not form joint bids under an FD policy, it is because the cost of bidding jointly is sufficiently high to outweigh the benefits of reduced competition. As a result, bidders' profits from solo bidding under FD are

5 Identification and Estimation

In my data, the winning firm observes all losing bids, i.e., an FD auction. In this section, I discuss how the parameters of the model in Section 4.1 are identified given the available data. I first introduce three additional empirically relevant features that were not included in the theoretical model. I then discuss the identification and estimation strategies.

5.1 Additional Empirical Features

First, in my data, a lease can produce both oil and gas. Therefore, I decompose the oil and gas quantities in each tract into an oil quantity, Q^o (measured in barrels), and a gas quantity, Q_g (measured in MCF, where 1 MCF equals 1,032 cubic feet). The total value of each tract is given by

$$Q = Q^o P^o + Q_g P^g,$$

where P^o and P^g are the prices of oil and gas, respectively. I choose P^o and P^g to be the annual average offshore U.S. Gulf Coast crude oil first purchase price and the Henry Hub spot price in the year of the auction, respectively.

Second, I introduce two additional cost components to account for development costs, which are costs that must be incurred after exploration if oil/gas is found to extract these resources. These additional cost components are $0 \leq \delta, \kappa \leq 1$, which are known to all bidders prior to the auction and are objects to be identified. δ represents yearly development costs such as maintenance and operation costs, and κ is an upfront development cost. Therefore, if the firm engages in exploratory drilling after winning the auction with a bid of b and observing that the exploration cost is $C = c$ and the production value is $Q = q$, its profit is $(\delta(1 - r) - \kappa)q - c - b$, where $\delta(1 - r) - \kappa$ is assumed to be positive.

Third, I assume that there exists a cost shifter $\mathcal{I} > 0$ observed by bidders prior to the auction such that $C = C^0 \zeta(\mathcal{I})$ where $C^0 > 0$ is a random variable with full support and with distribution $F^{C^0}(\cdot)$, and $\zeta : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is a differentiable and nonconstant function. Since it is not possible to distinguish the scale of $\zeta(\cdot)$ from the scale of C^0 , I also normalize $\zeta(\mathbb{E}(\mathcal{I})) = 1$. Conditional on $\mathcal{I} = \iota$, the upfront development cost is now $\kappa(\iota)$ where $\kappa : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is differentiable.

Assumption 2 C^0 , Q , and \mathcal{I} are pairwise independent.

higher than from joint bidding. Under an ND policy, the profit from solo bidding is even greater than under FD. Therefore, if the profit from joint bidding remains the same under both FD and ND—which occurs when all bidders form a single joint bid—bidders have even less incentive to form joint bids under ND.

Therefore, \mathcal{I} is a tract-level observed variable that affects the exploration and upfront development costs but does not carry any information about the tract's production value Q . C^0 represents the aspect of the exploration cost whose distribution F^{C^0} is common across all tracts. Henceforth, I refer to C^0 as the *base* exploration cost.

In my estimation, I construct an index that captures the lagged rig rental costs and use it as the cost shifter \mathcal{I} . Lagged rig rental costs are useful predictors of future rig rental costs, which are an important component of the exploration costs. To construct \mathcal{I} , I first examine the rig contracts that were signed within one year of and prior to a tract's auction date. Since a rig contract is signed between a lease operator (an oil/gas company) and a rig owner, for each rig contract, I identify all leases with the same operator and drilling activities during the contract's duration. The set of these leases are called matching leases. I then compute \mathcal{I} as the average rig rental costs of these prior contracts weighted by the distance between the tracts of matching leases to the tract being considered, accounting for area fixed effects and year fixed effects. Figure 27 in the Appendix shows the distribution of the constructed cost shifter across years and tract type. Because the cost shifter reflects the rig market condition at the tract's location before the auction occurs, a higher \mathcal{I} is intuitively associated with a lower exploration rate. Table 28 in the Appendix shows that this is indeed the case, providing evidence that this constructed cost shifter is relevant.

5.2 Identification

To facilitate exposition, I discuss identification while setting aside auction and tract heterogeneity. The objects of interest are $F^{\mathbf{S}}(\cdot)$, $F^{C^0}(\cdot)$, $V(\cdot)$, $\kappa(\cdot)$, $\zeta(\cdot)$ and δ . I observe bids, which directly identify the joint distribution of bids conditional on the cost shifter and royalty rate $F^{\mathbf{B}}(\cdot|\mathcal{I}, r)$, and drilling decision, which identify the probability of drilling $\Pr(d = 1|\mathbf{B}, \mathcal{I}, r)$.

First, it is without loss of generality to normalize the marginal distribution of S_i to be $U[0, 1]$ for all i .¹⁹ Let $F^B(\cdot)$ denote the marginal distribution of a firm's bid. Under this normalization, together with the symmetry assumption and that firms are playing $\beta^{FD}(\cdot)$, which is strictly increasing, firm i 's signal can be computed as follows

$$S_i = F^B(B_i|\mathcal{I}, r)$$

Therefore, $F^{\mathbf{S}}(\cdot)$ is identified. Using the first-order inversion of Guerre et al. (2000), $\Omega(s, \mathcal{I}, r) :=$

¹⁹By Sklar's theorem, by normalizing the signals to be uniformly distributed, the joint distribution that we identify is the copula of the true joint distribution.

$v(s, s|\mathcal{I}, r)$ is also identified.²⁰

Next, conditional on $\mathbf{S} = \mathbf{s}$, let $\Delta(\mathbf{s}, \mathcal{I}, r) := \Pr(d = 1|\mathbf{s}, \mathcal{I}, r)$ and $\mathcal{Z}(\mathcal{I}, r) := \frac{\delta(1-r) - \kappa(\mathcal{I})}{\zeta(\mathcal{I})}$. Therefore:

$$\Delta(\mathbf{s}, \mathcal{I}, r) = F^{C^0}(\mathcal{Z}(\mathcal{I}, r)V(\mathbf{s})) \quad (10)$$

Conditional on $S = s$ and $Y = s$, using integration by parts on $v(s, s|\mathcal{I}, r)$, it can be shown that:

$$\Omega(s, \mathcal{I}, r) = \begin{cases} \mathcal{Z}(\mathcal{I}, r)\zeta(\mathcal{I}) \int_0^s \int_0^{\tilde{z}} \Delta((s, s, \tilde{z}), \mathcal{I}, r) \frac{\partial}{\partial \tilde{z}} V((s, s, \tilde{z})) dz dG^{\mathbf{Z}}(z|s, s), & \text{if } N > 2 \\ \mathcal{Z}(\mathcal{I}, r)\zeta(\mathcal{I}) \int_0^s \Delta((s', s'), \mathcal{I}, r) \frac{\partial}{\partial s'} V((s', s')) ds', & \text{if } N = 2 \end{cases} \quad (11)$$

Proposition 3 $F^{C^0}(\cdot)$ and $\zeta(\cdot)$ are identified. $V(\cdot)$, $\kappa(\cdot)$, and δ are identified up to a multiplicative constant.


The proof of Proposition 3 is in Appendix C. Intuitively, $F^{C^0}(\cdot)$, $V(\cdot)$, $\zeta(\cdot)$, $\kappa(\cdot)$, and δ must satisfy both (10) and (11). The left-hand sides of both (10) and (11) are already identified from the data. Conditional on r , the slope of the level curves of $\Delta(\cdot, \cdot, r)$ is the same as the ratio of the derivative of the logs of $V(\cdot)$ and $\mathcal{Z}(\cdot, r)$. Because of the separability of $V(\cdot)$ and $\mathcal{Z}(\cdot, r)$, the logs of $V(\cdot)$ and $\mathcal{Z}(\cdot, r)$ are identified up to an affine transformation. They are further pinned down up to an additive constant using (11). Therefore, $V(\cdot)$ and $\mathcal{Z}(\cdot, r)$ are identified up to a multiplicative constant. From here, I identify F^{C^0} using the slope of $\Delta(\cdot, \mathcal{I}, r)$.

Figure 6 illustrates the identification argument outlined above, where \mathbf{s} is simplified to be unidimensional. The black lines in the figure represent the level curves of the expected drilling probability $\Delta(\cdot, \cdot, r)$ as a function of s and \mathcal{I} , and these are identified from the variation in drilling rates across (s, \mathcal{I}) . By fixing the value of \mathcal{I} , the ratio of the slopes of the level curves identifies the ratio of the derivative of $V(s)$. This implies that $\frac{V''(s)}{V'(s)}$ is identified.

²⁰Using the first-order inversion, conditional on $S = s$:

$$v(s, s|\mathcal{I}, r) = F^S(s|\mathcal{I}, r) + \frac{G(s|\mathcal{I}, r)}{g(s|\mathcal{I}, r)}$$

Figure 6: Illustration of the identification argument



With some algebra, it can be shown that:

$$\frac{\Delta_s''(s, \mathcal{I}, r)}{\Delta_s'(s, \mathcal{I}, r)} - \frac{V''(s)}{V'(s)} = \frac{F^{C^0''}(V(s)\mathcal{Z}(\mathcal{I}, r))}{F^{C^0'}(V(s)\mathcal{Z}(\mathcal{I}, r))}$$

Since the left-hand side of the equation is identified, F^{C^0} is the solution to a differential equation, with initial conditions ensuring that F^{C^0} remains bounded between 0 and 1. Intuitively, the identified level curves allow us to determine whether drilling rates are sensitive to changes in the signals. For instance, if drilling rates increase more significantly when signals improve for optimistic signals, it suggests either (1) that the cost distribution is convex, or (2) that the posterior mean function $V(\cdot)$ is convex. By fixing the cost shifter and comparing drilling rates at different signal levels, we can identify the concavity of $V(\cdot)$. Therefore, F^{C^0} is concave (or convex) when the second derivative of the level curves with respect to the signal (i.e., $\Delta_s''(\cdot, \mathcal{I}, r)$) is significantly more negative (or positive) than the rate of change of the slope of the level curves across different drilling rates (i.e., $\frac{V''(s)}{V'(s)}$).

Assumption 3 *The mean of the prior belief $\int_{supp(Q)} q f^Q dq$ is observed by the econometrician.*

Because the expected value of the posterior mean is always equal to the prior mean by Bayes' theorem, Assumption 3 provides one way to pin down the multiplicative constant in Proposition 3 by setting $\mathbb{E}(V(\cdot))$ to be equal to a known number. In my empirical application, I use Assumption 3 and set the prior mean to be $\mu_{Bernoulli, oil} \mu_{Gamma, oil} P^o + \mu_{Bernoulli, gas} \mu_{Gamma, gas} P^g$, using the variables constructed in Section 3. Further details are provided in the next section 5.3.

Corollary 1 *Under Assumption 3, $V(\cdot)$, $\kappa(\cdot)$, and δ are identified.*

Since $\mathcal{Z}(\mathcal{I}, r)$ and $\kappa(\cdot)$ are identified under Assumption 3, δ and $\zeta(\mathcal{I})$ can be separately identified using the variation in royalty rates in the data. During the data period, royalty rates vary between 12.5%, 16.67%, and 18.75%, depending on the year of the auction and the location of the tract being auctioned.

Much of my identification strategy relies on the variation in the drilling probability across bids and the presence of the cost shifter \mathcal{I} . This argument differs from the standard identification argument in the auction literature for two reasons. First, the standard method of identifying the joint distribution between the private signals and the latent value Q requires Q to be observed and utilizes the joint distribution of bids and Q .²¹ However, in my setting, Q is not observed when exploration does not take place. Second, for my counterfactual purposes, the primary objects that need to be identified are $V(\cdot)$ and $F^C(\cdot)$ (see, for example, the equilibrium bidding strategy $\beta^{ND}(\cdot)$). Even when Q is always observed, it is not possible to identify $F^C(\cdot)$ based only on the bid information without utilizing observed drilling decisions.

Discussion The model’s primitives do not include the prior belief F^Q and the signal generating process $F^S(\mathcal{S}|q)$ because all pairs of $\{F^Q, F^S(\mathcal{S}|Q)\}$ that generate the same $\{V(\cdot), F^S(\mathcal{S})\}$ are observationally equivalent in terms of both bidding and drilling behaviors. The most informative signal-generating process produces $Q = V(S)$, and the prior belief F^Q is the same as the distribution of $V(S)$. Any less informative signal generating process must be accompanied by a prior belief F^Q that is a mean-preserving spread of the distribution of $V(S)$.

5.3 Estimation

As discussed in Section 3, the primary source of observed heterogeneity between tracts arises from the spatial correlation in oil and gas deposits. For example, a tract located near a productive tract is more likely to be productive. This heterogeneity affects the distribution of the prior belief $F^Q(\cdot)$, which, in turn, influences the heterogeneity in $V(\cdot)$. As mentioned in Section 3, I assume that the heterogeneity across tracts due to spatial correlation in oil and gas deposits can be captured by the variables $(\mu_{\text{Bernoulli, oil}}, \mu_{\text{Bernoulli, gas}}, \mu_{\text{Gamma, oil}}, \mu_{\text{Gamma, gas}})$. Therefore, the object of interest is now $V(\cdot | \mu_{\text{Bernoulli, oil}}, \mu_{\text{Bernoulli, gas}}, \mu_{\text{Gamma, oil}}, \mu_{\text{Gamma, gas}})$.

²¹For a review of the literature, see Athey and Haile (2007)

However, due to the small sample size, it is not feasible to nonparametrically incorporate these variables into $V(\cdot)$. Thus, for estimation purposes, I parameterize $V(\cdot)$ as follows.

Let F^{Q^o} (F^{Q^g}) denote the distribution characterized by $\mu_{\text{Bernoulli, oil}}$ ($\mu_{\text{Bernoulli, gas}}$) and $\mu_{\text{Gamma, oil}}$ ($\mu_{\text{Gamma, gas}}$) as constructed in Section 3. The functional form of $V(\cdot)$ is micro-founded as the posterior mean given priors (F^{Q^o}, F^{Q^g}) and a Gaussian copula signal generating process.²² Specifically, for each tract, the joint distribution of $(S_1, S_2, \dots, S_N, Q^o, Q^g)$ is given by:

$$F^{\mathbf{S}, Q}(S_1, S_2, \dots, S_N, Q) = \Phi_R \left(\Phi^{-1} (F^S(S_1)), \dots, \Phi^{-1} (F^S(S_N)), \Phi^{-1} (F^{Q^o}(Q^o)), \Phi^{-1} (F^{Q^g}(Q^g)) \right)$$

where Φ_R is the joint CDF of a multivariate normal distribution with mean 0 and correlation matrix R and Φ^{-1} is the inverse of the CDF of a standard normal distribution. Since the bidders are ex ante symmetric, R must satisfy the following constraints:

$$\begin{cases} R_{ij} = R_{ik} & \text{for } i, j, k \leq N \\ R_{i, N+1} = R_{j, N+1} & \text{for } i, j \leq N \\ R_{i, N+2} = R_{j, N+2} & \text{for } i, j \leq N \end{cases}$$

R_{12} is thus the correlation between two bidders' private signals. To satisfy the assumption that the signals are affiliated, R_{12} must be weakly positive. $R_{1, N+1}$ and $R_{1, N+2}$ indicate the degree of correlation between the bidders' signals and the latent value of oil and gas, respectively.

The conditional distribution of $(\Phi^{-1} (F^{Q^o}(Q^o)), \Phi^{-1} (F^{Q^g}(Q^g))) | \mathbf{S}$ is thus

$$\mathcal{N} \left(\begin{bmatrix} M^o \sum_i S_i \\ M^g \sum_i S_i \end{bmatrix}, \Sigma_Q \right) \quad (12)$$

where

$$M^o = \frac{R_{1, N+1}}{R_{11} + R_{12}}, \quad M^g = \frac{R_{1, N+2}}{R_{11} + R_{12}}$$

$$\Sigma_Q = \begin{bmatrix} R_{N+1, N+1} & 0 \\ 0 & R_{N+2, N+2} \end{bmatrix} - \frac{N}{R_{11} + R_{12}} \begin{bmatrix} R_{1, N+1}^2 & R_{1, N+1} R_{1, N+2} \\ R_{1, N+1} R_{1, N+2} & R_{1, N+2}^2 \end{bmatrix}$$

²²As previously discussed, the resulting $V(\cdot)$ is also consistent with other priors that are mean-preserving contractions of F^{Q^o} and F^{Q^g} , provided they are accompanied by a less informative signal generating process.

The distribution specified in (12) then fully determines $V(\cdot)$. M^o and M^g quantify the marginal effects of the signals on the posterior belief about oil and gas quantities, respectively. Conditional on the prior belief $F^Q(\cdot)$, the larger M^o and M^g are, the more informative the signals. The precision of the distribution in (12) increases as the number of bidders increases. This represents the informational advantage of a winning bidder in an auction with more bidders because it receives more information ex post, thus having a more precise posterior belief.

My estimation procedure is conducted in two steps. In the first step, I use kernel estimation to estimate the CDF of the bids, which then yields the estimated private signals of the bidders. Since $V(\cdot)$ varies with N due to the informational advantage of having more bidders, the kernel estimation is performed on six separate samples, categorized by the number of bidders (2 bidders, 3 bidders, more than 3 bidders) and the tract's depth (shallow tracts (less than 400m deep) and deep tracts (at least 400m deep)).²³ For each subsample, I follow the method of Haile et al. (2003) to 'homogenize' the bids, accounting for auction-level observed heterogeneity, including year fixed effects, royalty rate, and tract water depth. To implement Haile et al. (2003)'s approach, I first regress the log of the observed bids on year fixed effects and a fifth-order spline of water depth for each subsample with the same royalty rate. The residuals are then treated as the log of the bids for auctions with homogeneous observed characteristics, conditional on \mathcal{I} and the four parameters that characterize the predicted production quantities ($\mu_{\text{Bernoulli, oil}}$, $\mu_{\text{Bernoulli, gas}}$, $\mu_{\text{Gamma, oil}}$, $\mu_{\text{Gamma, gas}}$). The homogenized bids preserve the optimality of the original bids when the observed heterogeneity commonly affects both the exploration cost and the tract's expected production value $V(\cdot)$ in a multiplicatively separable manner. I then estimate the marginal distribution of a firm's homogenized bid $F^B(B|\mathcal{I}, \mu_{\text{Bernoulli, oil}}, \mu_{\text{Bernoulli, gas}}, \mu_{\text{Gamma, oil}}, \mu_{\text{Gamma, gas}})$ using the Gaussian kernel and the least-squares cross-validation method to select bandwidths (Li et al., 2013). Given the small sample size, I simplify the set of conditioning variables to include only \mathcal{I} and the expected production value derived from the prior means ($\mu_{\text{Gamma, gas}}\mu_{\text{Bernoulli, gas}}P^g + \mu_{\text{Gamma, oil}}\mu_{\text{Bernoulli, oil}}P^o$). The estimated $F^B(\cdot)$ provides estimates of \mathcal{S} .

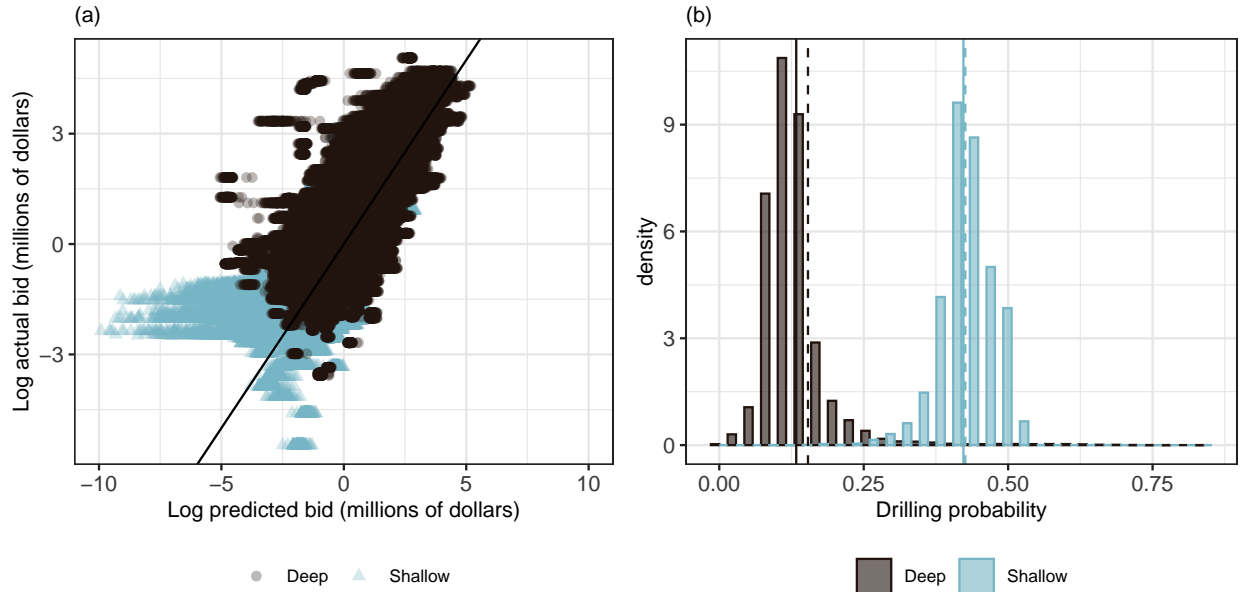
In the second step, the equilibrium bidding strategies $\beta^{FD}(\cdot)$ and the expected drilling probability before C is realized are computed for each draw of the model's parameters. I then estimate the parameters of the model by minimizing the sum of the squared difference between the predicted bids and the actual bids and the squared difference between the

²³The sample sizes of each subsample are listed in Table 25.

predicted drilling probabilities and the actual drilling outcomes. The standard errors are bootstrapped based on 100 iterations. To account for estimation errors in the predicted production quantities, each iteration incorporates a random draw from the estimated posterior distribution of the hyperparameters that define the Gaussian field used to generate these predicted production variables. Since the Gaussian field parameters are estimated using Bayesian inference with integrated nested Laplace approximations (Krainski et al., 2018), block bootstrapping was not required for this step.

5.4 Estimation results

Figure 7: Model fit



Note: The solid black line in panel (a) is the 45-degree line that goes through the origin. In panel (b), the solid lines are the mean of the fitted drilling probabilities, and the dashed lines are the average drilling rates in the data.

My model fit is summarized in Figure 7. In Panel (a), I plot the predicted bids, computed using the equilibrium bidding strategy under FD (Proposition 1), against the actual bids. The predicted median bids are \$367K and \$845K for shallow and deep tracts, which are close to the observed medians in the data of \$335K and \$850K. In Panel (b), I plot the distribution of the predicted drilling probability for deep and shallow tracts, which also fits well with the data. The average predicted drilling probabilities are 0.42 and 0.13 for shallow and deep tracts, similar to the average drilling rates of 0.43 and 0.15 in the data.

Table 8 summarizes the parameters that measure the informativeness of the bidders' private signals (M^o and M^g in (12)) and the correlation between these private signals (R_{12} in (12)). The correlation between bidders' signals is significantly nonnegative, thus implying that the bidders' signals are indeed affiliated. The estimated correlations are 0.11 for deep tracts and 0.12 for shallow tracts, which are significantly smaller than 1. These low correlations reflect the large differences between bids within the same auction, as observed in the data. Therefore, there is substantial uncertainty about other bidders' private information, which suggests that bid disclosure policies may be an effective means of influencing a bidder's posterior belief.

Table 8: Estimated parameters

Tract's depth	$Corr(S_i, S_j)$	Marginal effect of S on conditional mean of Q^o (M^o)	Marginal effect of S on conditional mean of Q^g (M^g)
Deep	0.11 (0.02)	0.94 (0.05)	0.01 (0)
Shallow	0.12 (0.04)	3.55 (2.14)	0.3 (0.01)


Note: $Corr(S_i, S_j)$ is the correlation between signals S_i in (12). The values in the second and third columns (M^o and M^g) are the increase in the expected value of oil and gas of the tract as S_i increases, measured as a percentile of the distribution of the posterior mean. Standard errors are bootstrapped.

The interpretations of M^o and M^g are as follows. Conditional on $S_i = 0.5 \forall i$, the posterior mean at the realized signal $V(\mathbf{S})$ is the same as the median of the prior belief. If one firm's signal increases to $S_i = 1$, the posterior mean conditional on the realized signals increases to the 68% percentile (51% percentile) of the prior distribution for oil (gas) for deep tracts. For shallow tracts, the increases are to the 96%-percentile (56%-percentile) of the prior distribution for oil (gas). My estimates imply that firms' private signals carry less information about gas than oil because most of the tracts explored during the data period are primarily for oil production.²⁴

Figure 9 shows the estimated shape of the CDF of the base drilling cost C^0 across deep and shallow tracts across all bootstrapped draws. The estimates of the parameters that characterize C^0 is in Table 33 in the Appendix. I make two observations. First, the estimated CDFs are significantly nonlinear. The p values of the Kolmogorov-Smirnov test

²⁴Much of the natural gas production in the Gulf of Mexico is from shallow tracts that were explored before my data period (Bureau of Ocean Energy Management, 2019).

Figure 9: Estimated shape of the CDF of the base drilling cost F^{C^0}



Note: Each line represents an estimated draw of the cost distribution F^{C^0} . The range of the cost estimates is similar to industry estimates. For example, the cost of a deep exploration well in the Gulf of Mexico in 2018 is said to be between 30 to 60 million on average (Sandrea and Stark, 2020), and prior to 2013 it was twice as expensive. For shallow exploration wells drilled after 1998, the average per-meter depth cost is about \$25,000/meter (Kaiser, 2021), which translates into less than \$10 million for shallow (less than 400m, as defined by the BOEM) wells.

against a linear function for F^{C^0} are less than 0.01 for both deep and shallow tracts. As explained in Section 4.2, a bid disclosure policy can affect the auctioneer's revenue only if the CDF of the exploration cost is nonlinear in the support of the posterior mean. Second, the CDF of the base exploration costs is weakly concave for deep tracts, but for shallow tracts, the cost distribution is neither concave nor convex.²⁵ Therefore, an ND policy is expected to weakly increase the drilling probability for deep tracts, but it is not theoretically clear whether an ND policy will improve the expected drilling probability for shallow tracts.

6 Counterfactual

In the counterfactual exercises, I study the effect of alternative bid disclosure policies on the bid revenue, drilling rate, and auctioneer revenue. I also consider the interaction between the royalty rate and disclosure policy. In these exercises, I fix the realizations of the bidders'

²⁵To test these hypotheses, I compute the 95% bootstrapped confidence interval of the minimum of the second derivative of F^{C^0} for each type of tract.

private signals as estimated from the data, and the royalty revenue is computed based on the expected value of the tract conditional on the realized signals.

I first discuss the effect of the ND policy in Section 6.1, followed by two partial disclosure policies in Section 6.2.1 and Section 6.2.2.

6.1 Nondisclosure Policy

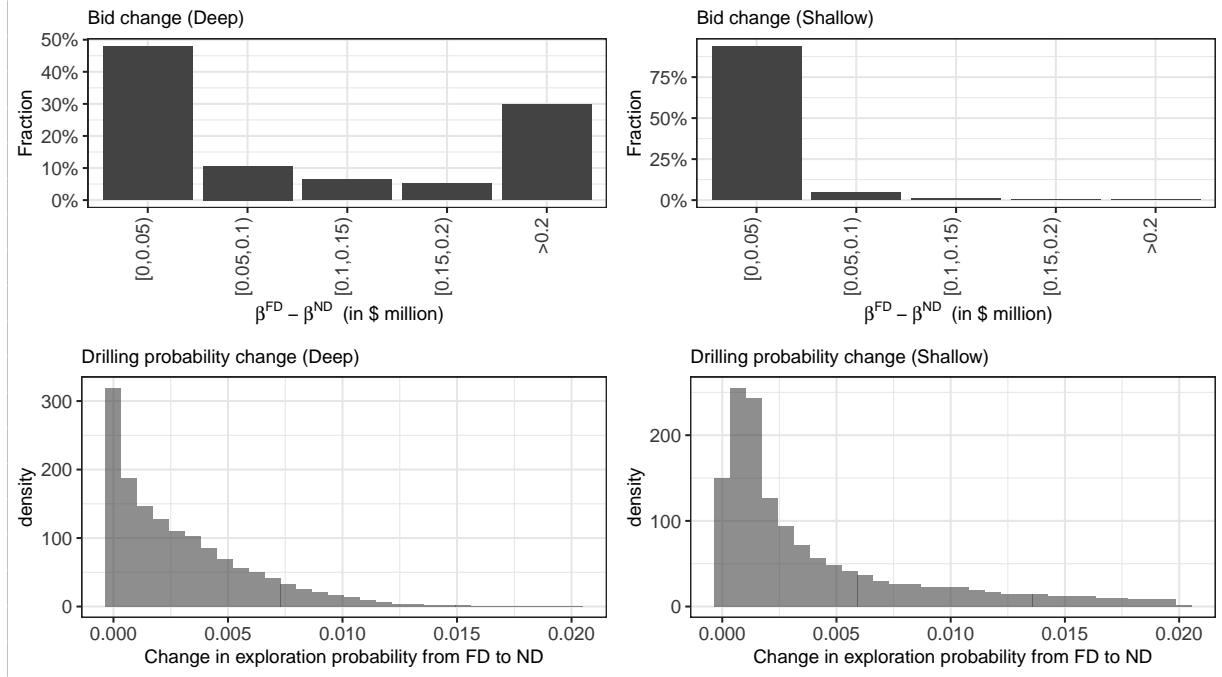
For each auction, I compare the drilling probability between the ND auction and the FD auction at the end of Stage 3, i.e., before the drilling cost is realized. Under ND, this drilling probability is determined when a bidder receives its own private signal S_i because it does not receive any further information ex post. Under FD, I compute the expected drilling probability before the losing bids are transmitted to the winning bidder. Throughout this analysis, I assume that the number of bidders is fixed. In Section 2 of the Online Appendix, I show that an ND policy improves revenue even when bidders' entry decisions are taken into account.

Figure 10 shows the distribution of the changes in the drilling probability (upper panels) and the bid revenue (lower panels) across tracts. As shown in Proposition 2, the ND policy results in lower equilibrium bids. For shallow tracts, the decrease is small for most tracts. For deep tracts, the distribution of the bid decrease is bimodal because of the large right tail of the bid distribution.²⁶ Since $F^{C^0}(\cdot)$ and thus $F^C(\cdot)$ is concave for deep tracts, the ND policy increases the drilling probabilities for all tracts. On the other hand, for shallow tracts, there is a small portion (0.85%) of tracts with a decrease in the drilling probability because $F^{C^0}(\cdot)$ (and thus $F^C(\cdot)$) is not everywhere concave for shallow tracts. Overall, the higher exploration rate under an ND policy results in an average royalty increase of \$0.99M per tract for deep tracts and \$0.82M for shallow tracts, significantly larger than the drop in bid revenue (Table 11). The large impact of the ND policy on the royalty revenue arises from the heterogeneity in the effects of the ND policy across tracts. Specifically, the increase in the drilling probability is higher for more productive tracts (Figure 30 in the Appendix).

Figure 12 demonstrates that the low correlation between bidders' signals plays a crucial role in the positive revenue gain resulting from the ND policy. In this exercise, I vary the correlation between bidders' signals, ranging from 0.1 to 0.9, while maintaining the same level of signal informativeness. As the signals become more correlated, the marginal value

²⁶The average bid for the tracts with a more than \$0.2M drop in bids has an average bid under FD of \$11.73M. As a comparison, the average bid under FD of tracts with between a \$0.15M to \$0.2M drop in bids is \$3.77M.

Figure 10: Distribution of the changes (per tract) in the drilling probability and the winning bids from FD to ND



Note: The changes in the drilling probability and the bid revenue are computed using the mean of the estimated parameters. The standard errors are not reported here. The bid revenue is computed at the realized signal of the highest bidder in the data and is in millions of dollars. The probability of drilling is between 0 and 1, and a positive change implies that an ND auction yields a higher expected drilling probability.

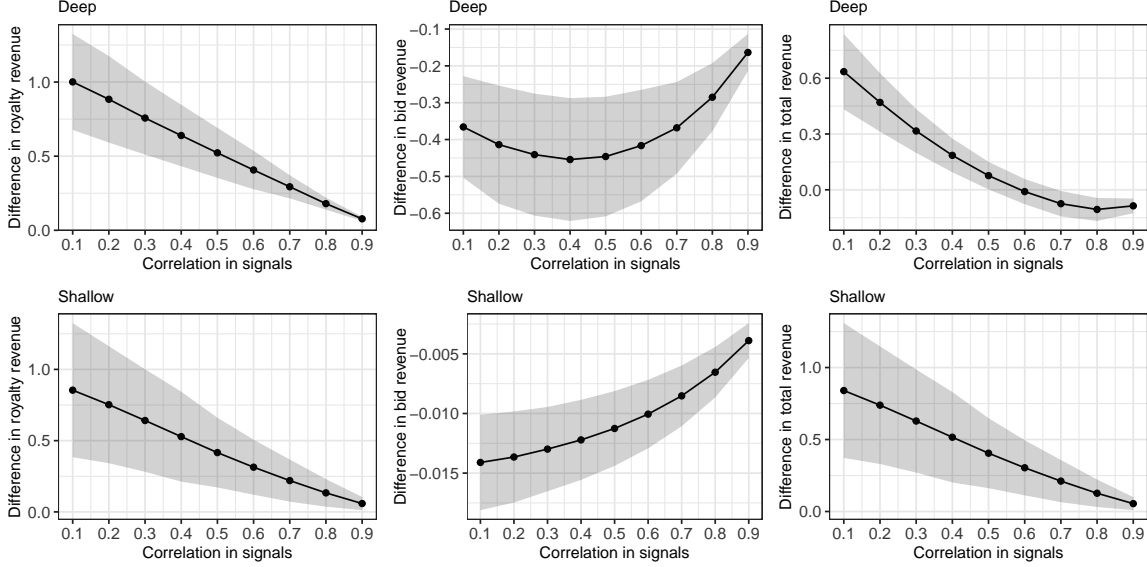
Table 11: Change in the bid and royalty revenues (per tract) from FD to ND (millions of dollars)

	Δ Bid	% Δ Bid	Δ Royalty (undiscounted)	Δ Royalty (discounted)*	% Δ Royalty
Deep	-0.37 (0.06)	-4.21 (0.33)	2.23 (0.4)	0.99 (0.18)	2.38 (0.25)
Shallow	-0.01 (0)	-1.13 (0.14)	1.84 (0.52)	0.82 (0.23)	1.61 (0.25)

Note*: To compute the discounted value, I assume a discount factor of 0.9 and that the royalty revenues are received in equal payments over 20 years. The numbers in the parentheses are bootstrapped standard errors.

of observing an additional signal (from a rival) diminishes. Consequently, the improvement in the probability of drilling under the ND policy decreases as signal correlation increases. Simultaneously, the bid revenue difference approaches zero as bidders possess relatively less private information, driving bids closer to the perfectly competitive benchmark. Overall, the

Figure 12: Revenue difference between ND and FD under different correlation values

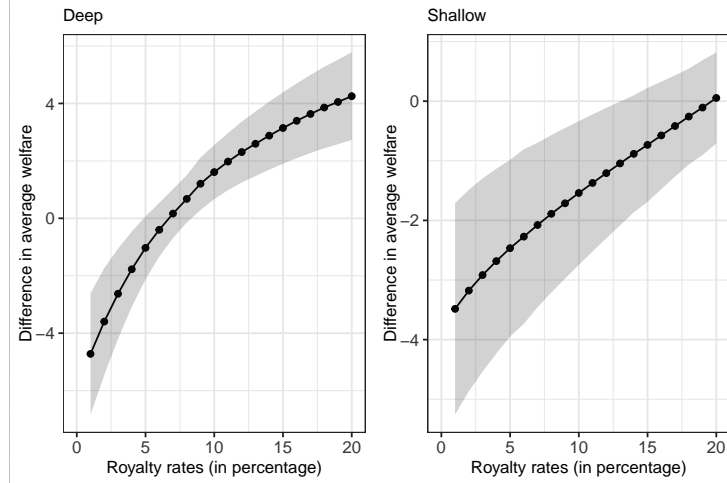


Note: Revenue is measured in millions of dollars. The grey band represents the bootstrapped 95% confidence interval.

change in total revenue under the ND policy becomes negative for deep tracts and is reduced by half for shallow tracts when the correlation exceeds 0.5.

Welfare discussion In addition to increasing the auctioneer’s revenue, the counterfactual estimates indicate that the ND policy does not reduce social welfare (Figure 13). Within the model, the FD policy creates an incentive misalignment between the auctioneer and the winning bidder due to the positive royalty rate. This is because the royalty rate functions as a tax on production *revenue*, discouraging the winning bidder from drilling at the socially optimal level. Consequently, the first-best drilling rate is defined as the drilling probability under the FD policy but in the absence of any royalty. In this analysis, welfare is measured as the mean absolute error between the first-best drilling probability for each tract and the drilling probability under a positive royalty rate, weighted by the production value of the tract, for a given disclosure policy (FD or ND). Figure 13 illustrates the welfare difference between FD and ND across various royalty rates r . The results show that at high royalty rates, the ND policy mitigates some of the distortions caused by the positive royalty rate, resulting in an overall welfare gain. At the current royalty rates observed in the data (approximately 12% across all tracts), the ND policy reduces the distortion caused by the positive royalty rate for 87.34% of deep tracts and 84.13% of shallow tracts. On average, the ND policy generates a \$2.84M in welfare gain for deep tracts and a smaller decrease in

Figure 13: Difference in welfare between a full-disclosure and a non-disclosure policy



Note: Welfare is measured as the mean absolute difference between the first-best drilling rate and the drilling rate under a disclosure policy (FD or ND) at a given royalty rate, weighted by the predicted production value of a tract (measured in millions of dollars). The y-axis in this figure represents the welfare difference between the ND and FD policies, where a positive value indicates a welfare gain under the ND policy. The shaded area represents the 95% bootstrapped confidence interval.

welfare (\$0.32M) for shallow tracts. Figure 34 in the Appendix provides the distribution of surplus changes across tracts at the current royalty rates.

6.2 Partial-disclosure Policy

Next, instead of considering the winner of the auction either observing all the submitted bids or not observing anything, I allow the auctioneer to commit to more general types of bid disclosure policies.

As explained in the ND analysis, switching away from the FD policy presents a tradeoff between the bid revenue and the royalty revenue. On the one hand, the ND policy decreases the bid revenue. On the other hand, the ND policy increases the expected royalty revenue. It is then natural to consider whether some intermediate policy, i.e., a partial disclosure policy, can perform better than the ND policy. A partial disclosure policy might outperform the ND policy for two reasons. First, the distribution of the drilling cost $F^C(\cdot)$ for shallow tracts is not concave; therefore, for shallow tracts, a partial disclosure policy might yield the highest expected drilling probability. Second, a partial disclosure policy might have a less adverse impact on the bid revenue than the ND policy.

In this section, I explore two partial disclosure policies that encompass both the FD and

ND policies. In the first exercise, I analyze a threshold disclosure policy that reveals losing bids only if they fall below a specified threshold. This policy is interesting for two reasons. First, it maximizes royalty revenue in the simple setting with binary drilling costs (*high* or *low*). This is because the auctioneer in this case can set the threshold such that the posterior mean when losing bids are undisclosed is equal to the *high* drilling cost, making the winning bidder indifferent between drilling and not drilling when the cost realization is high. Second, this threshold policy represents a special case of monotone partitional disclosure policies, which pool information across adjacent types. Such policies are commonly studied in other information disclosure settings and are optimal when the auctioneer’s revenue depends solely on the posterior mean (Dworczak and Martini, 2019). However, this optimality result does not extend to the OCS context, where the auctioneer’s ex post revenue is influenced not only by the posterior mean, which determines drilling probability, but also by the production value as well as by bid revenue.

In the second exercise, I examine a policy in which the auctioneer announces that the k highest bids, excluding the winning bid, will not be revealed. This policy has been implemented in other U.S. auction settings, such as failed financial institution auctions, where information about the second-highest bid is withheld ($k = 2$).²⁷ When a losing bid is not revealed, the winning bidder forms a belief about the hidden bid based on (1) the correlation between signals and (2) the values of the revealed bids, which constrain the possible range of the hidden bid. The decision to hide only the high losing bids ensures that the hidden bids are bounded above solely by the winning bid. Consequently, the value of the signals from these hidden bids is also bounded above by the winning bidder’s signal, which represents the most optimistic signal. As k increases, the weight of the winning bidder’s own information in shaping its posterior belief increases. In contrast, a policy that reveals high losing bids, by the same logic, increases the winning bidder’s sensitivity to the revealed losing bids. These revealed bids act as the upper bounds for the hidden bids, potentially reducing the expected drilling probability compared to the FD policy.

In addition, the two policies being considered are implementable in practice. In particular, they allow a bidder whose bid was only partially disclosed to verify that the auctioneer was truthful in their disclosed information. This ensures the auctioneer’s ability to credibly commit to a disclosure policy *ex ante*.²⁸

²⁷For more details on the bid disclosure policy in failed financial institution acquisition auctions, see <https://www.fdic.gov/resources/resolutions/bank-failures/failed-bank-list/biddocs.html>

²⁸On the other hand, policies that disclose only summary statistics (e.g., the mean) of losing bids cannot be independently verified by individual bidders and may be more challenging to implement.

6.2.1 Partial disclosure: Threshold disclosure

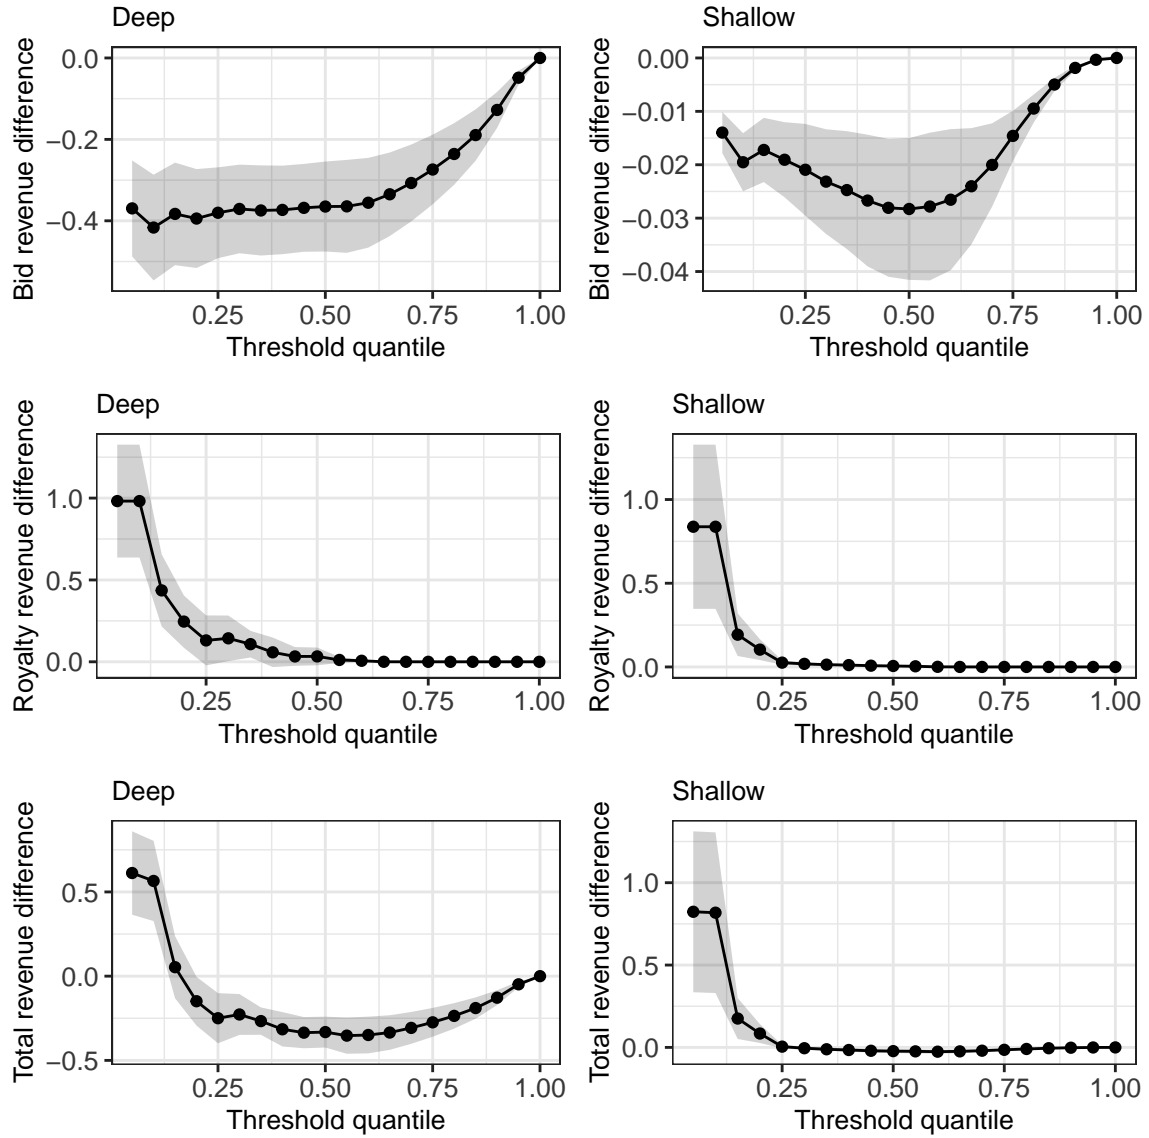
In this exercise, I consider the auctioneer’s revenue when the bid disclosure policy reveals the bids only if they are below a threshold α . When $\alpha = 0$, this is equivalent to the ND policy. When $\alpha = \infty$, this is the same as the FD policy. Proposition 3 in the Online Appendix establishes that a symmetric and increasing bidding strategy exists under a threshold disclosure policy.

To simplify the comparison between tracts, the threshold for each tract is normalized as follows. For a given value of $\alpha \in [0, 1]$, the threshold value of each tract is the α quantile value of the bid distribution of the same tract under the FD policy. For example, at $\alpha = 0.5$, the threshold of each tract is the median value of the equilibrium FD bid distribution of the tract. At $\alpha = 0$, this is now equivalent to the ND policy. At $\alpha = 1$, this is now equivalent to the FD policy.

Figure 14 illustrates the revenue composition for deep and shallow tracts under threshold disclosure policies at varying values of α . The results reveal a nonmonotonic relationship between the amount of information withheld (which decreases as the threshold α increases) and the bid revenue. For instance, for shallow tracts, the bid revenue when $\alpha = 0.5$ is lower than the bid revenue under ND. This occurs because a bidder whose signal exceeds α faces an additional trade-off when deciding its bid: by lowering its bid, the bidder gains access to more information post-auction but faces a lower winning probability. As $\alpha \rightarrow 1$ (the FD benchmark), this downward pressure on equilibrium bids diminishes since the benefit of additional information decreases. Similarly, as $\alpha \rightarrow 0$ (the ND benchmark), the loss from increased competition rises, and firms have less incentive to lower their bids. As a result, selecting an intermediate value of α may reduce revenue for the auctioneer. For example, setting $\alpha = 0.5$ decreases total revenue per tract by \$0.3645M for deep tracts and \$0.0278M for shallow tracts compared to the FD benchmark.

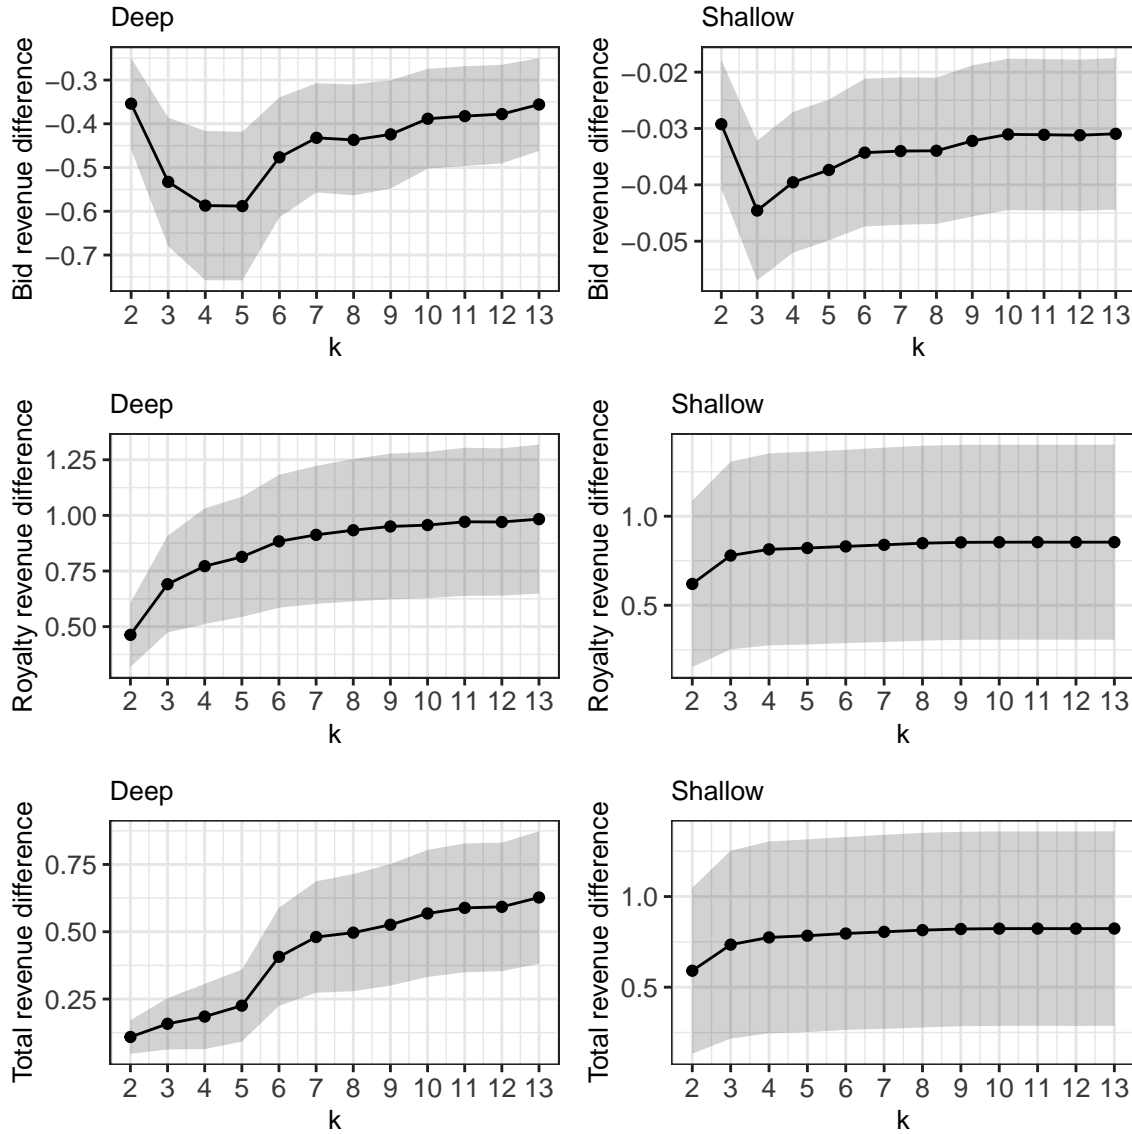
Figure 14 also shows that, on average, the ND policy generates the highest total revenue for both deep and shallow tracts. When considering disclosure policies tailored to individual tracts, 15.77% of deep tracts and 2.37% of shallow tracts achieve their optimal threshold at an intermediate value (Figure 36 in the Appendix). Additionally, the optimal threshold tends to be higher for tracts with greater oil and gas uncertainty but high potential (i.e., lower probability of finding oil or gas but high production value if discovered) (Table 37 in the Appendix).

Figure 14: Effect on the per-tract revenue of threshold disclosure policies relative to the FD benchmark (in millions of dollars)



Note: the x-axis of this graph is the α quantile of the equilibrium bid distribution under FD. For example, for $\alpha = 0.5$, the threshold being considered is the same as the median of the bid distribution for a particular tract under FD. The government's discount factor is assumed to be 0.9, and the royalty revenue for each tract is paid in equal payments over 20 years. The gray band is the 95% bootstrapped confidence interval.

Figure 15: Effect on the per-tract revenue of a withhold- k policy relative to the FD benchmark (in millions of dollars)



Note: When $k = 2$, only the second highest bid is not disclosed. When $k = 3$, the second and the third highest bid are not disclosed. The revenue on the y-axis is the revenue relative to the FD revenue. As $k \rightarrow N$, the bid revenue, the royalty revenue, and the total revenue all converge to those under the ND policy. In the computation of the royalty revenue, the government's discount factor is assumed to be 0.9, and the royalty revenue for each tract is paid in equal payments over 20 years. The gray band is the 95% bootstrapped confidence interval.

6.2.2 Partial disclosure: Withhold- k policy

In this exercise, I consider a partial disclosure policy in which the auctioneer announces that the k highest bids, not including the winning bid, are not revealed. For example, when $k = 4$,

the second, third, and fourth highest bids are not revealed. The ND policy is a special case in which $k = N$. Proposition 2 in the Online Appendix shows that the withhold- k policy has a symmetric and strictly increasing bidding strategy.

The revenue estimates of withhold- k policies are depicted in Figure 15. I make two observations. First, similar to the results from the threshold policy analysis, Figure 15 also exhibits nonmonotonicity in the bid revenue as the number of hidden losing bids increase. For example, for deep tracts, the policy that yields the lowest bid revenue is when $k = 5$. Second, on average, the withhold- k policy that yields the highest revenue for the auctioneer is still the ND policy. When the optimal withhold- k policy for each tract is considered (Figure 31 in the Appendix), the ND policy is optimal for almost all tracts. As a comparison, withholding only the second-highest bid achieves 17.3785% (71.7352%) of the revenue increase from the ND policy for deep (shallow) tracts.

Discussion In the model, I do not consider bidders' incentives to trade information after the auction under a partial or non-disclosure policy. Two institutional features support this assumption. First, the literature has noted that firms are reluctant to disclose their proprietary technologies for interpreting seismic surveys (Haile et al., 2010). For example, there have been documented cases where oil and gas companies refuse to share information with regulators due to concerns about losing their competitive advantage (Greene, 2011). Additionally, even information that must be submitted to regulators is subject to an embargo, remaining confidential for 50 years before being publicly released. Second, exploration by a winning bidder creates a positive externality for losing bidders who may be interested in bidding on adjacent tracts. In this scenario, the cost of revealing information for losing bidders is no longer negligible.

7 Conclusion

In this paper, I empirically study how information on the losing bids should be strategically revealed to the winning bidder in the context of US OCS auctions and quantify the potential revenue gain from such an information mechanism. In an OCS auction, the auctioneer (the government) auctions off the right to extract oil and gas on federal offshore tracts via a first-price sealed-bid auction. In addition to receiving the winning cash bid from the auction, the government also charges a royalty on the tract's production value. After the auction, the winning bidder decides whether to explore the tract, which is a costly decision. Unexplored

tracts do not produce oil and gas; therefore, the winning bidder’s post-auction action also affects the auctioneer’s payoff.

I first construct and estimate a model of a first-price sealed-bid pure common value auction in which the winner also chooses whether to explore the tract after the auction at a cost after observing all the losing bids. I show that by combining the bid data, the variation in the exploration rate across auctions, and an exploration cost shifter, the firms’ posterior beliefs on the production value of the tract conditional on all the bidders’ private signals are nonparametrically identified.

I then extend the model to the case in which the auctioneer can use an alternative bid disclosure policy. Under an alternative bid disclosure policy, the equilibrium bids decrease for two reasons. First, when less than full information is transmitted to the winning bidder, the winning bidder makes more ex post mistakes, which lowers the revenue from winning the auction, resulting in an incentive to lower its bid. Second, when the winning bidder is unable to observe the second-highest bid, the effect on the equilibrium bidding strategy is akin to a decrease in the competitiveness of the auction. On the other hand, an alternative bid disclosure policy might be able to improve the likelihood of exploratory drilling, which can ultimately result in higher royalty revenue. The counterfactual analysis reveals that an alternative bid disclosure policy can significantly improve the government’s revenue and offset some of the negative consequences of a positive royalty rate on welfare.

The economic trade-off discussed in this paper extends beyond OCS auctions to other contexts involving an auctioneer in a common-value setting, where the auctioneer’s utility depends on an ex post action taken by the winning bidder. For example, in pull supply contract auctions, where a buyer (the auctioneer) engages suppliers (the bidders), the winning supplier must invest in capacity before the buyer’s demand uncertainty is resolved. In this scenario, strategically disclosing losing bids could influence the winning bidder’s expectations about demand uncertainty, thereby mitigating the risk of underinvestment in capacity. A similar rationale applies to timber and failed bank auctions, where bidders can gain insights into bidder-specific persistent private cost components—such as capacity constraint in timber auctions or liquidity constraint in failed bank auctions—based on revealed bids. This transparency could dampen competition and reduce the auctioneer’s revenue in future auctions.

References

- ALLEN, J., R. CLARK, B. R. HICKMAN, AND E. RICHERT (2022): “Resolving Failed Banks: Uncertainty, Multiple Bidding & Auction Design,” .
- ATHEY, S. AND P. A. HAILE (2007): “Nonparametric approaches to auctions,” *Handbook of econometrics*, 6, 3847–3965.
- ATHEY, S. AND J. LEVIN (2001): “Information and competition in US forest service timber auctions,” *Journal of Political economy*, 109, 375–417.
- BAJARI, P., S. HOUGHTON, AND S. TADELIS (2014): “Bidding for incomplete contracts: An empirical analysis of adaptation costs,” *American Economic Review*, 104, 1288–1319.
- 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, 1319–1363.
- BLONDEL, V. D., J.-L. GUILLAUME, R. LAMBIOTTE, AND E. LEFEBVRE (2008): “Fast unfolding of communities in large networks,” *Journal of statistical mechanics: theory and experiment*, 2008, P10008.
- BOEM (2019): “Exploration and Development Scenario: Lease Sale 258,” Accessed: 2023-10-01.
- BUREAU OF OCEAN ENERGY MANAGEMENT (2019): “OCS Report, BOEM 2021-005,” Tech. rep.
- DEBROCK, L. M. AND J. L. SMITH (1983): “Joint bidding, information pooling, and the performance of petroleum lease auctions,” *The Bell Journal of Economics*, 395–404.
- DWORCZAK, P. AND G. MARTINI (2019): “The simple economics of optimal persuasion,” *Journal of Political Economy*, 127, 1993–2048.
- ESÓ, P. AND B. SZENTES (2007): “Optimal information disclosure in auctions and the handicap auction,” *The Review of Economic Studies*, 74, 705–731.
- FAJT, M. (2009): “Bids for failed banks: a state secret?” *American Banker*.
- GREENE, M. A. (2011): “Spilling Secrets: Trade Secret Disclosure and Takings in Offshore Drilling Regulation,” *Richmond Journal of Law & Technology*, 17.

- GUERRE, E., I. PERRIGNE, AND Q. VUONG (2000): “Optimal nonparametric estimation of first-price auctions,” *Econometrica*, 68, 525–574.
- HAILE, P., K. HENDRICKS, AND R. PORTER (2010): “Recent US offshore oil and gas lease bidding: A progress report,” *International Journal of Industrial Organization*, 28, 390–396.
- HAILE, P. A., H. HONG, AND M. SHUM (2003): “Nonparametric tests for common values at first-price sealed-bid auctions,” Tech. rep., National Bureau of Economic Research.
- HENDRICKS, K., J. PINSKE, AND R. H. PORTER (2003): “Empirical Implications of Equilibrium Bidding in First-Price, Symmetric, Common Value Auctions.” *Review of Economic Studies*, 70.
- HENDRICKS, K. AND R. H. PORTER (1992): “Joint bidding in federal OCS auctions,” *The American Economic Review*, 82, 506–511.
- (1996): “The timing and incidence of exploratory drilling on offshore wildcat tracts,” *The American Economic Review*, 388–407.
- (2014): “Auctioning resource rights,” *Annu. Rev. Resour. Econ.*, 6, 175–190.
- HODGSON, C. (2018): “Information Externalities, Free Riding, and Optimal Exploration in the UK Oil Industry,” .
- JOSEPH, S. (2020): “Why Google’s decision to withhold programmatic data is pushing some advertisers to pull ad spending,” *Digiday*.
- KAISER, M. J. (2021): “A Review of exploration, development, and production cost offshore Newfoundland,” *Natural Resources Research*, 30, 1253–1290.
- KAISER, M. J., B. SNYDER, AND A. G. PULSIPHER (2013): “Offshore drilling industry and rig construction market in the Gulf of Mexico,” Tech. rep., United States. Bureau of Ocean Energy Management. Gulf of Mexico OCS Region.
- KAMENICA, E. AND M. GENTZKOW (2011): “Bayesian persuasion,” *American Economic Review*, 101, 2590–2615.
- KONG, Y. (2021): “Sequential auctions with synergy and affiliation across auctions,” *Journal of Political Economy*, 129, 148–181.

- KRAINSKI, E. T., V. GÓMEZ-RUBIO, H. BAKKA, A. LENZI, D. CASTRO-CAMILO, D. SIMPSON, F. LINDGREN, AND H. RUE (2018): *Advanced spatial modeling with stochastic partial differential equations using R and INLA*, CRC Press.
- KRASNOKUTSKAYA, E., K. SONG, AND X. TANG (2020): “The role of quality in internet service markets,” *Journal of Political Economy*, 128, 75–117.
- KRISHNA, V. AND J. MORGAN (1997): “(Anti-) competitive effects of joint bidding and bidder restrictions,” *Working paper*.
- LEWIS, G. AND P. BAJARI (2011): “Procurement contracting with time incentives: Theory and evidence,” *The Quarterly Journal of Economics*, 126, 1173–1211.
- LI, Q., J. LIN, AND J. S. RACINE (2013): “Optimal bandwidth selection for nonparametric conditional distribution and quantile functions,” *Journal of Business & Economic Statistics*, 31, 57–65.
- LIN, C.-Y. C. (2009): “Estimating strategic interactions in petroleum exploration,” *Energy Economics*, 31, 586–594.
- (2013): “Strategic decision-making with information and extraction externalities: A structural model of the multistage investment timing game in offshore petroleum production,” *Review of Economics and Statistics*, 95, 1601–1621.
- MILGROM, P. R. AND R. J. WEBER (1982): “A Theory of Auctions and Competitive Bidding,” *Econometrica*, 50, 1089–1122.
- PORTER, R. H. (1995): “The Role of Information in US Offshore Oil and Gas Lease Auction,” *Econometrica: Journal of the Econometric Society*, 1–27.
- RUE, H., S. MARTINO, AND N. CHOPIN (2009): “Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations,” *Journal of the royal statistical society: Series b (statistical methodology)*, 71, 319–392.
- SANDREA, R. AND P. STARK (2020): “Deep Water Oil Exploration in the Gulf of Mexico: A Spatial Model Provides Clues for Undiscovered Potential,” .
- SOMAINI, P. (2020): “Identification in auction models with interdependent costs,” *Journal of Political Economy*, 128, 3820–3871.

TAKAHASHI, H. (2018): “Strategic design under uncertain evaluations: structural analysis of design-build auctions,” *The RAND Journal of Economics*, 49, 594–618.

VREUGDENHIL, N. (2020): “Booms, busts, and mismatch in capital markets: Evidence from the offshore oil and gas industry,” *Journal of Political Economy*, *Forthcoming*.

XIAO, M. AND Z. YUAN (Forthcoming): “License complementarity and package bidding: The us spectrum auctions,” *American Economic Journal: Microeconomics*.

A Proof of Proposition 1

Proof:

Fix any strictly increasing $\beta(\cdot)$ and suppose that all firms $j \neq i$ play according to β . Firm i 's expected profit from bidding b when $S_i = s$ is given by Eq (3). Differentiating Eq (3) with respect to b , the FOC is

$$0 = -G^Y(\beta^{-1}(b)|s) + \frac{1}{\beta'(\beta^{-1}(b))} [v(s, \beta^{-1}(b)) - b] g^Y(\beta^{-1}(b)|s)$$

At a symmetric equilibrium in which every firm plays β , the above equation must be satisfied at $b = \beta(s)$, which implies that

$$\begin{aligned} 0 &= -G^Y(s|s) + 1/\beta'(s) [v(s, s) - \beta(s)] g^Y(s|s) \\ \iff \beta'(s) + \beta(s) \frac{g^Y(s|s)}{G^Y(s|s)} &= v(s, s) \frac{g^Y(s|s)}{G^Y(s|s)} \end{aligned} \quad (13)$$

Eq (13) is a linear ordinary differential equation with a boundary condition of $\beta(0) = 0$. Therefore, Eq (3) is the unique solution of the FOC.

Next, I show that $\beta(\cdot)$ is increasing. Because the signals are affiliated, $v(s, y)$ is increasing in both s and y . Moreover, for $\tilde{s} > s$, $L(y|s) < L(y|\tilde{s})$ for any y . Therefore, $L(\cdot|s)$ is first order stochastic dominated by $L(\cdot|\tilde{s})$. Therefore, $\beta(\cdot)$ is indeed an increasing function.

Finally, I show that $\beta(\cdot)$ is firm i 's best response against all other firms playing β . Suppose that firm i bids $\beta(\tilde{s})$, where $\tilde{s} \neq s$. Its expected profit is

$$U(\tilde{s}|s) := \int_0^{\tilde{s}} [v(s, y) - \beta(\tilde{s})] g^Y(y|s) dy$$

Take the derivative with respect to \tilde{s} :

$$\frac{\partial U(\tilde{s}|s)}{\partial \tilde{s}} = \left([v(s, \tilde{s}) - \beta(\tilde{s})] \frac{g^Y(\tilde{s}|s)}{G^Y(\tilde{s}|s)} - \beta'(\tilde{s}) \right) G^Y(\tilde{s}|s) \quad (14)$$

From the FOC, $\frac{\partial U(\tilde{s}|s)}{\partial \tilde{s}} = 0$. Eq (14) can be rewritten as

$$\frac{\partial U(\tilde{s}|s)}{\partial \tilde{s}} = \left([v(s, \tilde{s}) - \beta(\tilde{s})] \frac{g^Y(\tilde{s}|s)}{G^Y(\tilde{s}|s)} - [v(\tilde{s}, \tilde{s}) - \beta(\tilde{s})] \frac{g^Y(\tilde{s}|\tilde{s})}{G^Y(\tilde{s}|\tilde{s})} \right) G^Y(\tilde{s}|s) \quad (15)$$

Because $v(s, \tilde{s}) < v(\tilde{s}, \tilde{s})$ and $\frac{g^Y(\tilde{s}|s)}{G^Y(\tilde{s}|s)} < \frac{g^Y(\tilde{s}|\tilde{s})}{G^Y(\tilde{s}|\tilde{s})}$ by the affiliation property, the LHS of Eq (15) is negative. Therefore, bidding $\beta(\tilde{s})$ is not firm i 's best response.

Similarly, if $\tilde{s} < s$, then $v(s, \tilde{s}) > v(\tilde{s}, \tilde{s})$ and $\frac{g^Y(\tilde{s}|s)}{G^Y(\tilde{s}|s)} > \frac{g^Y(\tilde{s}|\tilde{s})}{G^Y(\tilde{s}|\tilde{s})}$. Therefore, bidding $\beta(\tilde{s})$ is also not firm i 's best response. ■

B Proof of Proposition 2

Proof:

Fix any strictly increasing $\beta(\cdot)$ and suppose that all firms $j \neq i$ play according to β . After the auction, conditional on winning at a bid b , firm i 's profit before the exploration cost is realized is given by

$$\hat{V}_0(s, \beta^{-1}(b)) = \mathbb{E} [V(s, y, \mathbf{z}|s, y \leq \beta^{-1}(b))]$$

where we have defined \hat{V}_0 on Eq (7).

Therefore, firm i 's expected profit from bidding b when $S_i = s$ is given by

$$\int_0^{\beta^{-1}(b)} \left(\psi \left(\hat{V}_0(s, \beta^{-1}(b)) \right) - b \right) g^Y(y|s) dy \quad (16)$$

The FOC w.r.t b is

$$0 = \left[\psi \left(\hat{V}_0(s, \beta^{-1}(b)) \right) - b \right] g^Y(\beta^{-1}(b)|s) \frac{1}{\beta'(\beta^{-1}(b))} + \left(\frac{\partial}{\partial b} \psi \left(\hat{V}_0(s, \beta^{-1}(b)) \right) - 1 \right) G^Y(\beta^{-1}(b)|s) \quad (17)$$

where

$$\begin{aligned}
&= \frac{\partial}{\partial b} \left[1/G^Y(\beta^{-1}(b)) \int^{\beta^{-1}(b)} \int^y V(s, y, \mathbf{z}) g^Z(\mathbf{z}|s, y) g^Y(y|s) d\mathbf{z} dy \right] \\
&= \left(\bar{V}_0(s, \beta^{-1}(b)) - \hat{V}_0(s, \beta^{-1}(b)) \right) \frac{g^Y(\beta^{-1}(b)|s)}{\beta'(\beta^{-1}(b)) G^Y(\beta^{-1}(b)|s)}
\end{aligned} \tag{18}$$

At a symmetric equilibrium in which every firm plays β , the above equation must be satisfied at $b = \beta(s)$.

$$\begin{aligned}
0 &= \left[\psi \left(\hat{V}_0(s, s) \right) - \beta(s) \right] g^Y(s|s) \frac{1}{\beta'(s)} + \\
&\quad \left(\psi' \left(\hat{V}_0(s, s) \right) \left[\bar{V}_0(s, s) - \hat{V}_0(s, s) \right] \frac{g^Y(s|s)}{\beta'(s) G^Y(s|s)} - 1 \right) G^Y(s|s) \\
&\implies \beta'(s) + \beta(s) \frac{g^Y(s|s)}{G^Y(s|s)} = \left[\psi \left(\hat{V}_0(s, s) \right) + \psi' \left(\hat{V}_0(s, s) \right) \left[\bar{V}_0(s, s) - \hat{V}_0(s, s) \right] \right] \frac{g^Y(s|s)}{G^Y(s|s)}
\end{aligned} \tag{19}$$

Equation (20) is a linear ODE equation with an initial condition $\beta(s) = 0$. Therefore, Eq 9 is the solution of (20).

It remains to show that $\psi \left(\hat{V}_0(s, s) \right) + \psi' \left(\hat{V}_0(s, s) \right) \left[\bar{V}_0(s, s) - \hat{V}_0(s, s) \right]$ is increasing in s . Note that $\psi'(\cdot) = F^C(\cdot)$, therefore, $\psi(\cdot)$ is convex. It follows that, for $s' > s$:

$$\begin{aligned}
&\psi \left(\hat{V}_0(s', s') \right) - \psi \left(\hat{V}_0(s, s) \right) \\
&\geq \psi' \left(\hat{V}_0(s, s) \right) \left[\hat{V}_0(s', s') - \hat{V}_0(s, s) \right] \\
&= \psi' \left(\hat{V}_0(s, s) \right) \left[\hat{V}_0(s', s') - \bar{V}_0(s, s) \right] - \psi' \left(\hat{V}_0(s, s) \right) \left[\bar{V}_0(s, s) - \hat{V}_0(s, s) \right] \\
&\geq \psi' \left(\hat{V}_0(s, s) \right) \left[\hat{V}_0(s', s') - \bar{V}_0(s', s') \right] - \psi' \left(\hat{V}_0(s, s) \right) \left[\bar{V}_0(s, s) - \hat{V}_0(s, s) \right]
\end{aligned}$$

where the last inequality follows from $\bar{V}_0(s, s) < \bar{V}_0(s', s')$. Therefore, $\beta(\cdot)$ is increasing.

The rest of the proof (showing that $\beta(\cdot)$ is a bidder's best response) follows the same logic as in Proof A and is therefore omitted. ■

C Proof of Proposition 3

Proof:

As stated in the main text, $F^{C^0}(\cdot)$, $V(\cdot)$, $\zeta(\cdot)$, $\kappa(\cdot)$, and δ are objects to be identified, and they must satisfy both (10) and (11). The left-hand side (LHS) of (10) and (11) are identified from the data.

Consider (10). The slopes of the level curves of $\Delta(\cdot, \cdot, r)$ are also identified and equal to $-\frac{\frac{\partial \log \Delta(\mathbf{s}, \mathcal{I}, r)}{\partial \mathbf{s}}}{\frac{\partial \log \Delta(\mathbf{s}, \mathcal{I}, r)}{\partial \mathcal{I}}}$. Therefore:

$$\frac{\frac{\partial \log \Delta(\mathbf{s}, \mathcal{I}, r)}{\partial \mathbf{s}}}{\frac{\partial \log \Delta(\mathbf{s}, \mathcal{I}, r)}{\partial \mathcal{I}}} = \frac{\frac{\partial \log V(\mathbf{s})}{\partial \mathbf{s}}}{\frac{\partial \log \mathcal{Z}(\mathcal{I}, r)}{\partial \mathcal{I}}} \quad (21)$$

Suppose that there exists $\{\tilde{V}(\cdot), \tilde{\zeta}(\cdot), \tilde{F}^{C^0}, \tilde{\zeta}(\mathcal{I}), \tilde{\delta}, \tilde{\kappa}\}$ and $\{V(\cdot), \zeta(\cdot), F^{C^0}, \zeta(\mathcal{I}), \delta, \kappa\}$ that both satisfy (10), (11), and thus (21). From (21), for any \mathbf{s}, \mathcal{I} , and r . Similar to the main text, I define

$$\mathcal{Z}(\mathcal{I}, r) := \frac{\delta(1-r) - \kappa(\mathcal{I})}{\zeta(\mathcal{I})}$$

$$\tilde{\mathcal{Z}}(\mathcal{I}, r) := \frac{\tilde{\delta}(1-r) - \tilde{\kappa}(\mathcal{I})}{\tilde{\zeta}(\mathcal{I})}$$

Thus,

$$\frac{\frac{\partial}{\partial \mathbf{s}} \log \tilde{V}(\mathbf{s})}{\frac{\partial}{\partial \mathbf{s}} \log V(\mathbf{s})} = \frac{\frac{\partial}{\partial \mathcal{I}} \log \tilde{\mathcal{Z}}(\mathcal{I}, r)}{\frac{\partial}{\partial \mathcal{I}} \log \mathcal{Z}(\mathcal{I}, r)}$$

Since the LHS of the above equation is a function of only \mathbf{s} whereas the right-hand side (RHS) is a function of only (\mathcal{I}, r) , there exists two constants $k_0, k_1, k_2 > 0$ such that

$$\tilde{\mathcal{Z}}(\mathcal{I}, r) = k_2 \mathcal{Z}(\mathcal{I}, r)^{k_0}$$

and $\tilde{V}(\mathbf{s}) = k_1 V(\mathbf{s})^{k_0}$

WLOG, let $k_0 \geq 1$. I now prove that $k_0 = 1$. Suppose, for a contradiction, that $k_0 > 1$. Consider two cases:

Case 1: $N > 2$. From Eq (11), conditional on $S_i = s$

$$\begin{aligned}
0 &= \int^s \int^z \Delta(s, s, \tilde{\mathbf{z}}, I) \left[\mathcal{Z}(\mathcal{I}, r) \zeta(\mathcal{I}) \frac{\partial}{\partial \tilde{\mathbf{z}}} V(s, s, \tilde{\mathbf{z}}) - \tilde{\mathcal{Z}}(\mathcal{I}, r) \tilde{\zeta}(\mathcal{I}) k_1 \frac{\partial}{\partial \tilde{\mathbf{z}}} V^{k_0}(s, s, \tilde{\mathbf{z}}) \right] d\tilde{\mathbf{z}} dG^{\mathbf{Z}}(\mathbf{z}|s, s) \\
&= \int^s \int^z \Delta(s, s, \tilde{\mathbf{z}}, I) \left[\frac{\partial}{\partial \tilde{\mathbf{z}}} V(s, s, \tilde{\mathbf{z}}) \left(\mathcal{Z}(\mathcal{I}, r) \zeta(\mathcal{I}) - \tilde{\mathcal{Z}}(\mathcal{I}, r) \tilde{\zeta}(\mathcal{I}) k_1 k_0 V^{k_0-1}(s, s, \tilde{\mathbf{z}}) \right) \right] d\tilde{\mathbf{z}} dG^{\mathbf{Z}}(\mathbf{z}|s, s) \\
&\geq \int^s \int^z \Delta(s, s, \tilde{\mathbf{z}}, I) \left[\frac{\partial}{\partial \tilde{\mathbf{z}}} V(s, s, \tilde{\mathbf{z}}) \left(\mathcal{Z}(\mathcal{I}, r) \zeta(\mathcal{I}) - k_1 k_0 V^{k_0-1}(s, s, \tilde{\mathbf{z}}) \right) \right] d\tilde{\mathbf{z}} dG^{\mathbf{Z}}(\mathbf{z}|s, s)
\end{aligned} \tag{22}$$

where the last inequality is due to $\tilde{\mathcal{Z}}(\mathcal{I}, r) \tilde{\zeta}(\mathcal{I}) = \tilde{\delta}(1-r) - \tilde{\kappa}(\mathcal{I}) \leq 1$. Since $V(\cdot)$ is strictly increasing, $\frac{\partial}{\partial \tilde{\mathbf{z}}} V(s, s, \tilde{\mathbf{z}}) > 0$. Since $V(\mathbf{0}) = 0$, for $S_i = s$ sufficiently small, the right hand side of Eq. (22) is positive, which is a contradiction.

Case 2: $N = 2$. From (11), conditional on $S_i = s$

$$0 = \int^s \Delta(s', s', I) \left(\mathcal{Z}(\mathcal{I}, r) \zeta(\mathcal{I}) \frac{\partial}{\partial s'} V(s', s') - \tilde{\mathcal{Z}}(\mathcal{I}, r) \tilde{\zeta}(\mathcal{I}) k_1 \frac{\partial}{\partial s'} V^{k_0}(s', s') \right) ds' \tag{23}$$

Similar to the $N > 2$ case, for $S_i = s$ sufficiently small, the right hand side of Eq. (23) is positive, which is a contradiction.

Therefore, $k_0 = 1$. Therefore, $V(\cdot)$ and $\mathcal{Z}(\cdot, \cdot)$ are identified up to two multiplicative constant k_1 and k_2 , respectively.

I now show that $\zeta(\cdot)$, $\kappa(\cdot)$, and δ are also identified up to a multiplicative constant. Note that if $\mathcal{Z}(\mathcal{I}, r)$ is identified, $\delta/\zeta(\mathcal{I})$ and $\frac{\kappa(\mathcal{I})}{\zeta(\mathcal{I})}$ is identified due to the variations in r in the data. Using the normalization $\zeta(\mathbb{E}(\mathcal{I})) = 1$, δ is identified. Therefore, $\zeta(\mathcal{I})$ and $\kappa(\mathcal{I})$ are both identified. In other words

$$\begin{aligned}
\tilde{V}(\mathbf{s}) &= k_1 V(\mathbf{s}) \\
\tilde{\zeta}(\mathcal{I}) &= \zeta(\mathcal{I}) \\
\tilde{\kappa}(\mathcal{I}) &= k_2 \kappa(\mathcal{I}) \\
\tilde{\delta}(\mathcal{I}) &= k_2 \delta
\end{aligned} \tag{24}$$

Since (11) must be satisfied, $k_2 = \frac{1}{k_1}$. Therefore, $V(\cdot)$, $\kappa(\cdot)$, and δ are identified up to the same multiplicative constant, and $\zeta(\cdot)$ is identified.

I now show that F^{C^0} is also identified as long as $V(\cdot)$, $\kappa(\cdot)$, and δ is identified. Consider (10)

and take derivative of $\Delta(\cdot, \mathcal{I}, r)$ at \mathbf{s} and $\mathbf{s}' \neq \mathbf{s}$

$$\frac{\partial/\partial \mathbf{s} \Delta(\mathbf{s}, \mathcal{I}, r)}{\partial/\partial \mathbf{s} \Delta(\mathbf{s}', \mathcal{I}, r)} = \frac{f^{C^0}(V(\mathbf{s}) \mathcal{Z}(\mathcal{I}, r)) V'(\mathbf{s})}{f^{C^0}(V(\mathbf{s}') \mathcal{Z}(\mathcal{I}, r)) V'(\mathbf{s}')}$$

and

$$\frac{\partial/\partial \mathbf{s} \Delta(\mathbf{s}, \mathcal{I}, r)}{\partial/\partial \mathbf{s} \Delta(\mathbf{s}', \mathcal{I}, r)} = \frac{\tilde{f}^{C^0}(\tilde{V}(\mathbf{s})(\mathcal{I}, r)) \tilde{V}'(\mathbf{s})}{\tilde{f}^{C^0}(\tilde{V}(\mathbf{s}')(\mathcal{I}, r)) \tilde{V}'(\mathbf{s}')}$$

The above equations are well defined because $V(\cdot)$ is strictly increasing. Using (24) and that $k_2 = 1/k_1$, the above two equations imply

$$\frac{f^{C^0}(V(\mathbf{s}) \mathcal{Z}(\mathcal{I}, r))}{\tilde{f}^{C^0}(V(\mathbf{s}) \mathcal{Z}(\mathcal{I}, r))} = \frac{f^{C^0}(V(\mathbf{s}') \mathcal{Z}(\mathcal{I}, r))}{\tilde{f}^{C^0}(V(\mathbf{s}') \mathcal{Z}(\mathcal{I}, r))} \quad (25)$$


This implies that there exists a constant k_3 such that

$$\tilde{f}^{C^0}(\cdot) = k_3 f^{C^0}(\cdot) \quad (26)$$

Since f^{C^0} and \tilde{f}^{C^0} are both density functions, $k_3 = 1$. Therefore, f^{C^0} and thus F^{C^0} are identified. ■


D Tables and Figures

Figure 16: Average exploration rate, tract depth, and hit rate over time



Note: The left panel and middle panel show the average exploration rate and average water depth among tracts that were assigned within a given year. The right panel shows the proportion of tracts that were successfully developed among tracts that were explored.

Figure 17: Number of explored leases among newly acquired leases for each oil/gas company in 3-year rolling windows



Note: Each dot in the graph represents the number of newly acquired leases and explored leases within a 3-year window for a particular year for each oil/gas company. The number of leases includes only leases with at least 50% ownership. The three-year benchmark was chosen because, during my data period, most explored leases (68.5851%) were explored within the first three years. Most firms acquire many more leases than the number of tracts that will eventually be explored, and some firms do not explore any leases. For example, on average, firms acquire 28 leases but explore only 2 in a three-year window.

Table 18: Correlation between the number of explored leases and the number of all newly acquired leases for each owner

	Number of explored leases
Number of newly added leases (at least 50% ownership)	0.029*** (0.008)
Year FE	Yes
Lease owner FE	Yes
Observations	633
Adjusted R ²	0.543

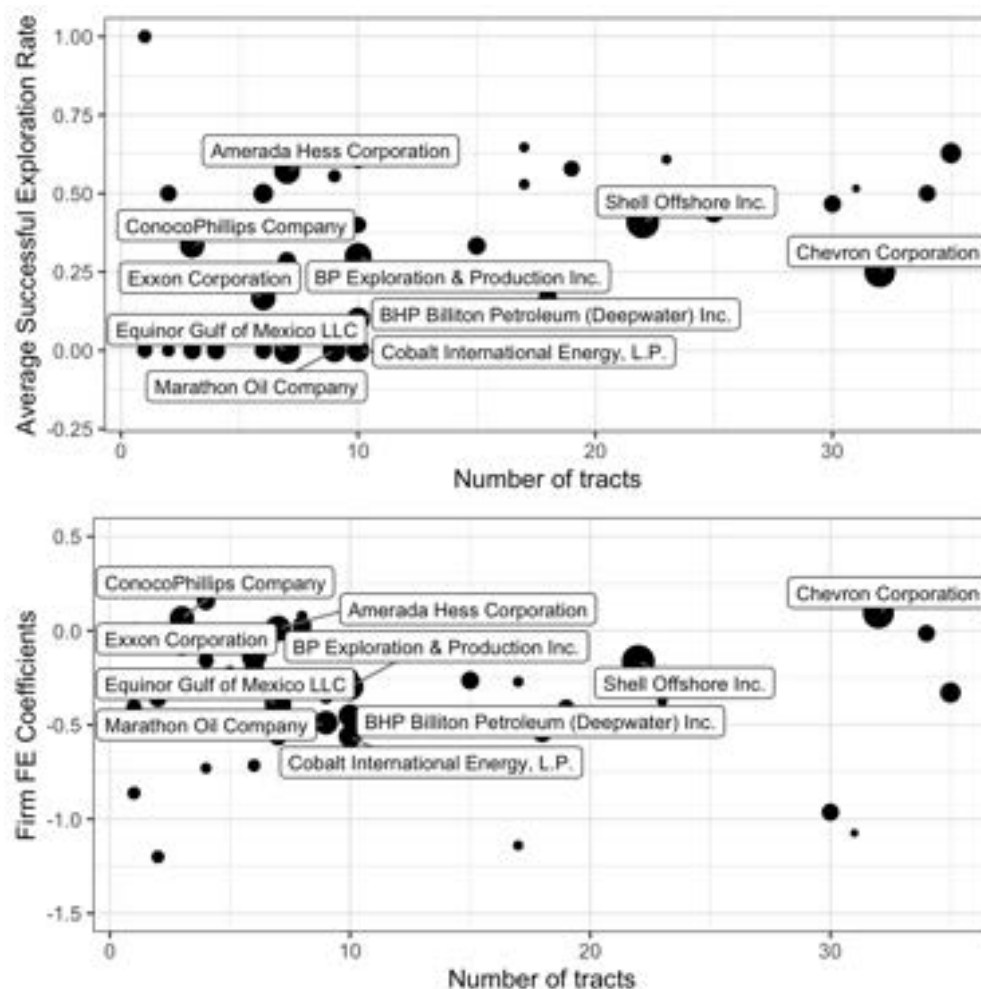
Note: In this table, I regress the number of explored leases among newly acquired tracts in a given year for owners with majority ownership percentage ($\geq 50\%$) on the total number of tracts acquired in the same year, year fixed effects, and owner fixed effects. Standard errors are in parentheses and adjusted for heteroskedasticity. The positive correlation suggests that gas and oil companies are more likely to explore new tracts if they have more newly acquired tracts. Therefore, exploration capacity is unlikely to be the main reason for low exploration rates.

Table 19: Summary statistics of all available auctions and production data

Statistic	N	Mean	St. Dev.	Max	Min
First bid	16,280	1,463	5,194	157,111	0.12
Second bid	46	1,555	5,015	84,391	4.24
Average losing bid	46	16	2,823	56,964	4.24
Number of bids	16,280	0	0	0.01	0
Fraction of Explored Tracts	30,757	0.28	0.45	1	0
Oil production	4,119	4,702	19,446	581,873	0
Gas production	4,119	40,098	109,538	4,253,872	0

Note: The bids are expressed in hundreds of thousands of dollars. Oil production is measured in thousand barrels of crude oil, and gas production is measured in million cubic feet.

Figure 20: Successful exploration rate across large oil/gas companies



Note: This figure includes only oil and gas companies classified as large bidders in Section F, i.e., those whose total bids exceed that of the median bidder during the data period. The horizontal axis represents the number of explored tracts owned by each company, including those acquired through joint bids. In the top figure, the vertical axis shows the average success rate among explored tracts. In the bottom figure, the vertical axis represents the coefficients of firm fixed effects from a regression where I estimate whether a tract was successfully explored. The regression includes the log value of the winning bid, the log value of the average losing bid, the number of bidders, the log value of the tract's water depth, the predicted quantities of oil and gas in the tract, as well as fixed effects for the tract's location and year. The size of each point in the figure reflects the total bid amount placed by each firm during the data period.

Table 21: Summary statistics of the total number of geological and geophysical survey permit application by lessees

(a) All permits

	Fraction of tracts with survey permits	Number of tracts
Unexplored tracts	0.21	22094
Explored unproductive tracts	0.17	4544
Explored productive tracts	0.28	4119

(b) Only 3D and 4D permits

	Fraction of tracts with survey permits	Number of tracts
Unexplored tracts	0.10	22094
Explored unproductive tracts	0.08	4544
Explored productive tracts	0.15	4119

Note: The number of permits indicates whether a tract's lessee applied for a geological and geophysical permit after acquiring the tract but before any exploration drilling occurs. A permit is required prior to surveying; however, the location of the survey is proprietary information and is not disclosed by the government. Thus, this value represents the upper bound of the actual number of tract-specific permits. In the upper table, all permits are included. In the lower table, only deep-penetration seismic 3D and 4D surveys are included, as these are typically the most informative for exploration and development.

Figure 22: Current composition of source of revenue


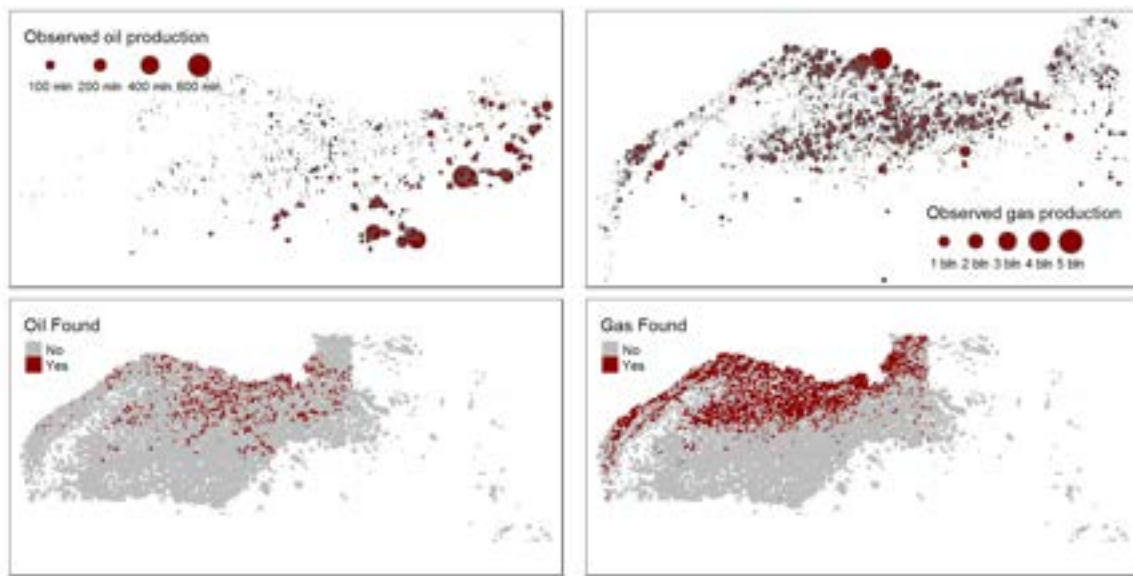
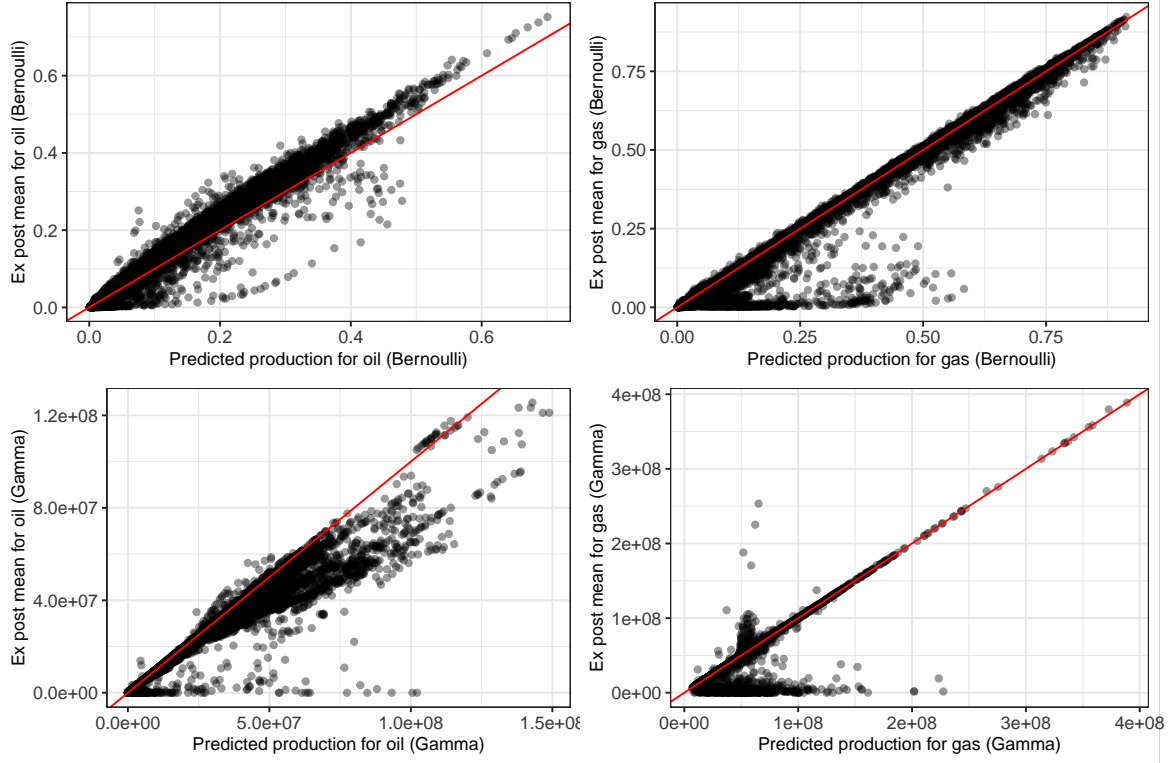


Figure 23: Observed oil and gas production



Note: The upper graphs show the locations and production quantities of tracts that have produced oil or gas or both. The units of production are barrels for oil and MCF for gas. The lower graphs show the locations of tracts that were explored. The gray tracts are tracts that were explored but were subsequently abandoned, and the red tracts are tracts that were explored and subsequently developed. The lower graphs cover a larger area than the upper graphs.

Figure 24: Comparisons between ex post and prior constructed beliefs on oil and gas



Note: The red line represents the 45-degree line. The x-axis shows the mean of the prior distribution, while the y-axis shows the mean of the ex post distribution. The unit for the Gamma mean is in MCF.

Table 25: Subsample statistics

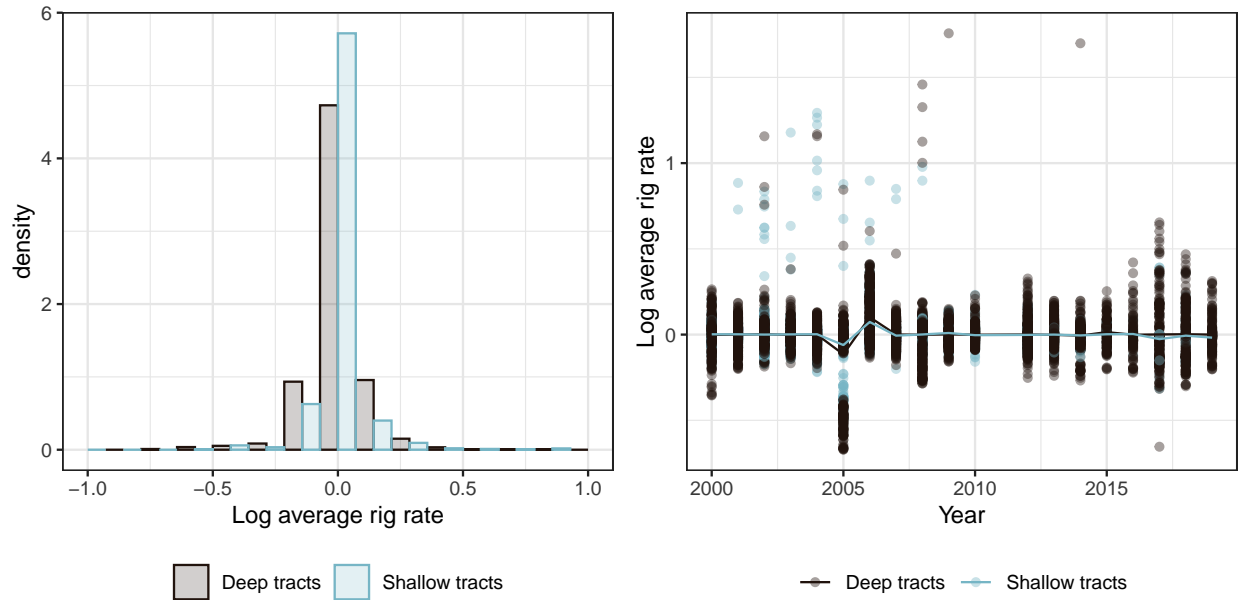
Number of bidders	Tract type	Percentage	Number of observations
2	Deep	9%	1025
2	Shallow	3%	343
3	Deep	52%	5714
3	Shallow	5%	555
> 3	Deep	2%	186
> 3	Shallow	29%	3127

Table 26: Summary statistics of rig contracts

Statistic	Mean	N	St. Dev.
Contract Duration (Days)	143.50	4,649	275.50
Daily Rate	90.82	4,649	113.70
Total Rig Cost	27,551.00	4,649	121,522.00

Note: The rig contracts are only available between 2000 and 2019. The data is recorded at the contract level. The units of the daily rate and the total rig rental cost are thousands of dollars.

Figure 27: Distribution of cost shifter across years and tract type




Note: The solid lines in the right panel represent the average annual values of the constructed cost shifter over the years.

Table 28: Relevance of Cost Shifter

	Explored		
	(1)	(2)	(3)
Log Average Rig Rate	-0.1425 (0.0545)	-0.1077 (0.0544)	-0.1474 (0.0541)
Number of Bids		0.0294 (0.0078)	0.0146 (0.0079)
Log First Bid		0.0363 (0.0102)	0.0564 (0.0100)
Log Average Losing Bid		0.0229 (0.0137)	0.0494 (0.0137)
Area FE	No	Yes	Yes
Time FE	No	No	Yes
Observations	2,510	2,510	2,510
Adjusted R ²	0.0015	0.1623	0.2475

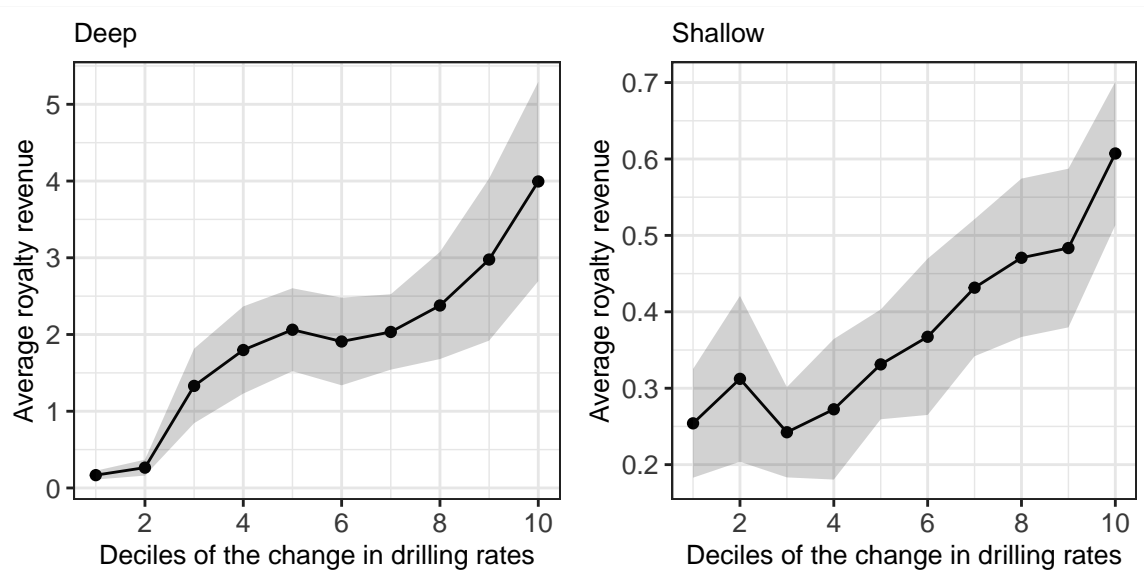
Note: The rig rental costs are in thousands of dollars. The average rig rental cost is the constructed cost shifter \mathcal{Z} , which is a tract-level variable. It represents the distance-weighted daily rig rental costs of nearby tracts prior to a tract's auction, accounting for year fixed effects and area fixed effects. The regression estimates suggest that this constructed average daily rig rental cost is a useful predictor of the realized cost of exploration, as it is negatively correlated with the exploration probability.

Figure 29: Average and standard deviation of contract rates across operators



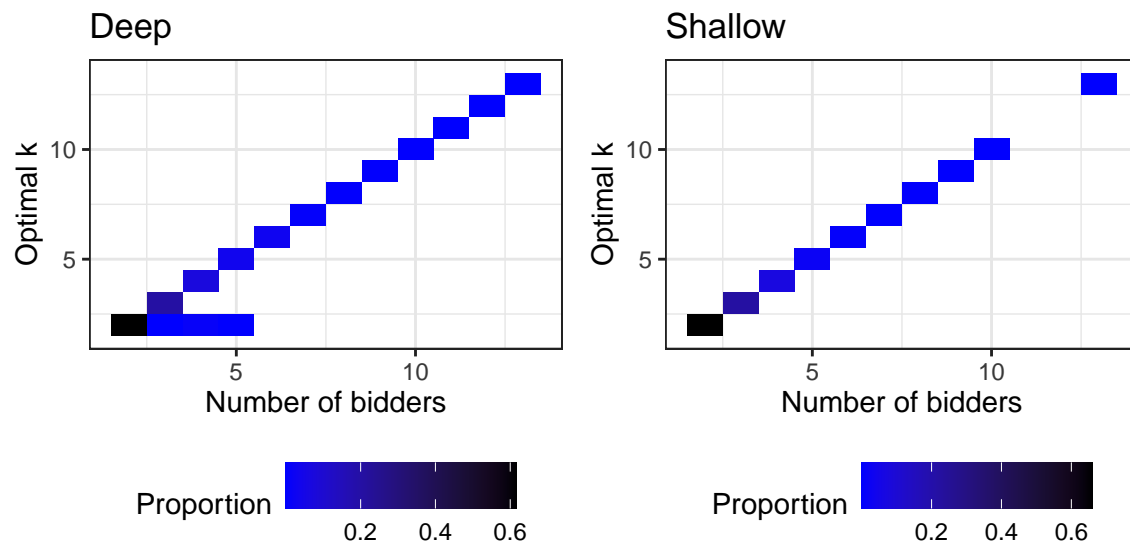
Note: Each contract is signed between a rig owner and an oil/gas company, which acts as the operator of a lease. The contract outlines the daily price for rig rental. On the graph, each point represents a summary statistic of all rig contracts signed by the same operator, with the mean shown on the left and the standard deviation on the right.

Figure 30: Average royalty revenue (in millions of dollars) across deciles of drilling rate change



Note: In this figure, I categorized tracts into 10 groups according to the impact of a non-disclosure policy on the tracts' drilling probabilities. For example, a tract in the 5th decile has a median increase in the drilling probability relative to other tracts in the sample. On the vertical axis, I compute the average royalty revenue of those tracts. This figure shows that the most productive deep tracts tend to be associated with the largest increase in the drilling probability when a non-disclosure policy is implemented. This explains the significant increase in the royalty revenue for deep tracts despite the modest average increase in the drilling probability.

Figure 31: Distribution of optimal k values across tracts



Note: In this exercise, I find the optimal withhold- k policy for each tract. When $N = k$, the optimal policy is the same as the ND policy.

Table 32: Correlation between losing bids and exploration rate

	Exploration drilling = TRUE			
	(1)	(2)	(3)	(4)
Winning bid (in log)	0.047 (0.010)	0.047 (0.010)	0.048 (0.010)	0.047 (0.010)
Average losing bid (in log)	0.039 (0.012)	0.039 (0.012)	0.038 (0.012)	0.028 (0.013)
Number of bids	0.008 (0.008)	0.006 (0.008)	0.007 (0.008)	0.007 (0.009)
Predicted production for oil (Bernoulli)		-0.049 (0.134)	-0.503 (0.452)	-0.306 (0.553)
Predicted production for gas (Bernoulli)		0.258 (0.122)	0.614 (0.288)	0.526 (0.315)
Predicted production for oil (Gamma) (in log)		0.032 (0.018)	0.031 (0.018)	0.015 (0.019)
Predicted production for gas (Gamma) (in log)		-0.022 (0.025)	-0.020 (0.025)	-0.007 (0.027)
Ex post mean for oil (Bernoulli)			0.422 (0.411)	0.262 (0.512)
Ex post mean for gas (Bernoulli)			-0.395 (0.318)	-0.213 (0.349)
Observations	2,510	2,510	2,510	1,988
Adjusted R ²	0.315	0.316	0.316	0.369

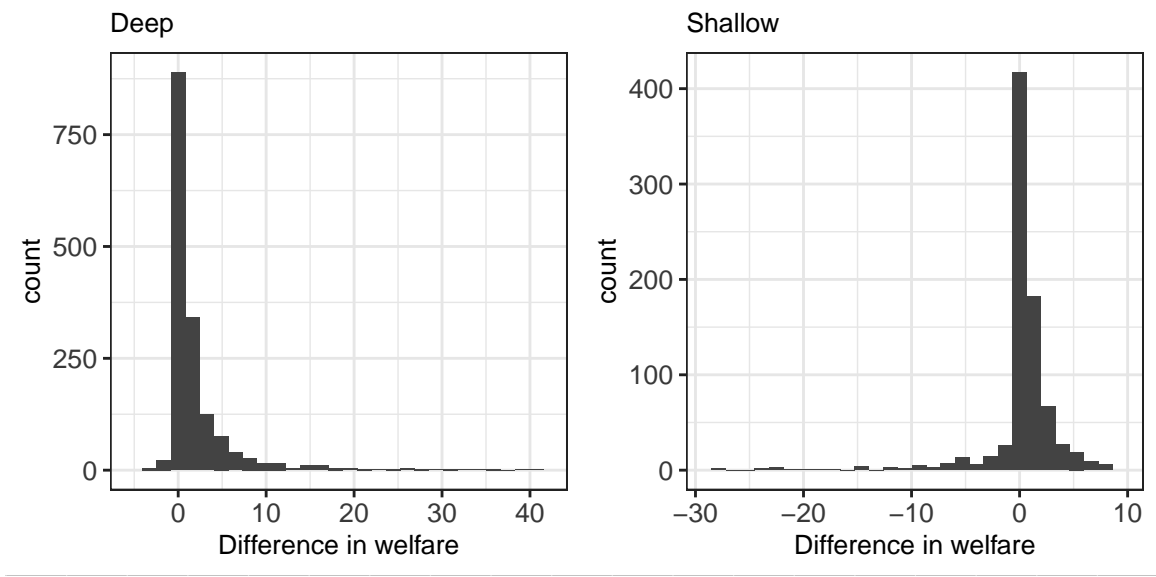
Note: Other included but unreported variables are the number of tracts owned by lessees that have not been explored or are under production, those that have been explored, and those currently under production. Additionally, the model includes the average bids from neighboring tracts, the number of neighboring tracts explored, the number under production, tract depth, year fixed effects, and area fixed effects. In Column (4), leases with shared ownership of nearby tracts are excluded. Standard errors are in parentheses and adjusted for heteroskedasticity.

Table 33: Estimated parameters related to F^{C^0}

Deep	-2.47 (0.053)	-4.91 (0.087)	-4.44 (0.084)	-3.87 (0.067)	-3.92 (0.029)	-0.55 (0.067)	-1.62 (0.006)	-2.3 (0.001)	-1.28 (0)
Shallow	-0.94 (0.042)	-4.01 (0.032)	-4.07 (0.05)	-3.73 (0.054)	-3.19 (0.089)	-1.84 (0)	-1.86 (0)	-1.86 (0)	-0.83 (0)

Note: For each tract depth (shallow or deep), F^{C^0} is modeled using a cubic spline function with five interior knots. These knots are positioned approximately at the 10th, 30th, 50th, 70th, and 90th percentiles of the estimated posterior mean distribution (i.e., the distribution of $V(\mathcal{S})$). Bootstrapped standard errors are provided in parentheses.

Figure 34: Distribution of welfare difference across tracts at observed royalty rates



Note: Surplus is measured in millions of dollars. A negative (positive) surplus indicates that the FD policy results in a higher (lower) surplus compared to the ND policy.

Table 35: Correlation between the exploration decision and the values of losing bids

	Exploration drilling = TRUE			
	(1)	(2)	(3)	(4)
Winning bid (in log)	0.059 (0.010)	0.058 (0.010)	0.047 (0.009)	0.047 (0.009)
Second highest bid (in log)	0.044 (0.012)	0.044 (0.013)	0.036 (0.011)	0.037 (0.011)
Difference between third highest bid and winning bid (in log)	-0.029 (0.010)	-0.022 (0.012)	-0.021 (0.009)	-0.017 (0.011)
At most two bids	-0.391 (0.139)	-0.294 (0.166)	-0.264 (0.124)	-0.217 (0.148)
Difference between fourth highest bid and winning bid (in log)		-0.018 (0.017)		-0.008 (0.016)
At most three bids		-0.274 (0.244)		-0.111 (0.223)
Ex post mean variables included	No	No	Yes	Yes
Observations	2,510	2,510	2,510	2,510
Adjusted R ²	0.318	0.318	0.474	0.473

Note: Other control variables, not reported here, are consistent with those in Column (5) of Table 3. These include the number of tracts owned by lessees that are unexplored, under production, or have been explored; the average bids from neighboring tracts; the number of neighboring tracts explored or under production; tract depth; year and area fixed effects; the number of bids; and predicted production variables. Ex post belief variables are defined as in Table 3. Standard errors, presented in parentheses, are adjusted for heteroskedasticity.

Figure 36: Distribution of optimal thresholds (under the threshold disclosure policy) across tracts

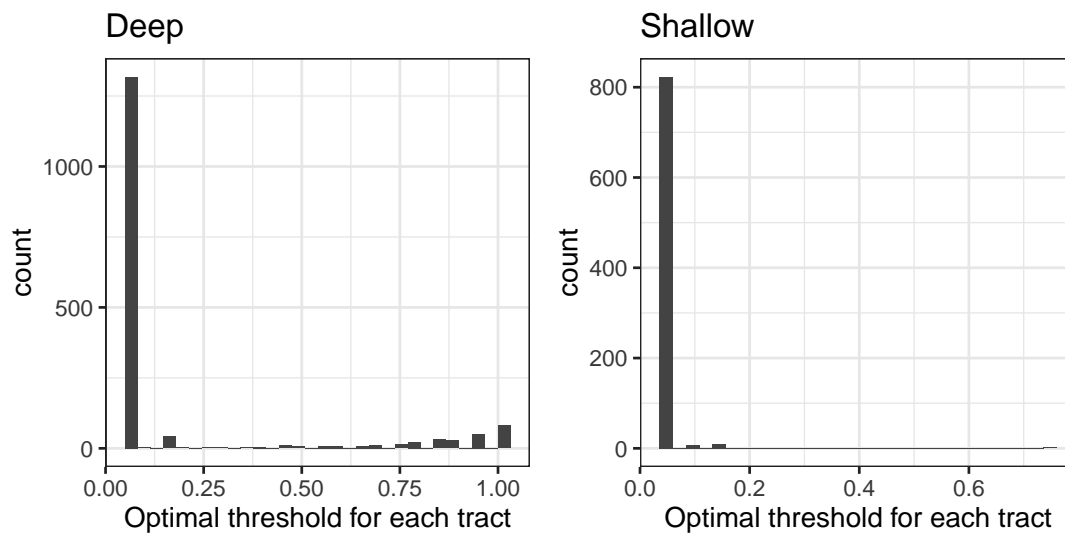


Table 37: Correlation between the optimal threshold and expected production quantities

	Optimal threshold for each tract	
	Deep	Shallow
	(1)	(2)
Predicted production for oil (Gamma) (in log)	0.407 (0.155)	0.095 (0.068)
Predicted production for gas (Gamma) (in log)	0.596 (0.356)	0.177 (0.091)
Predicted production for oil (Bernoulli)	-15.260 (1.246)	-0.700 (0.386)
Predicted production for gas (Bernoulli)	0.326 (1.504)	-0.860 (0.502)
Observations	1,661	843

E Construction of tract’s potential ($\mu_{gamma,oil}$, $\mu_{gamma,gas}$, $\mu_{bernoulli,oil}$, $\mu_{bernoulli,gas}$)

I assume that once production begins at a tract, all potential bidders observe that tract’s production value. For leases that have expired—which I term ‘completed’ leases—the oil and gas quantities of a tract are equal to the total amount of oil and gas the tract has produced. For tracts that are still in production at the end of my data period, I construct a measure of the total amount of oil and gas that will be produced using the production data of the previous 2 years.

Construction of production values For leases that have not expired by the end date of my data set, I use the following regression on completed leases to predict the total amount of oil and lease in the tract:

$$\log \left(\sum_{\tau=T}^{\infty} Q_t^{i,\tau} \right) = \beta_0 + \sum_{j=1}^{\underline{T}} \beta_j Q_t^{i,T-j} + \alpha X_t + \gamma_T + \epsilon_t^i \quad (27)$$

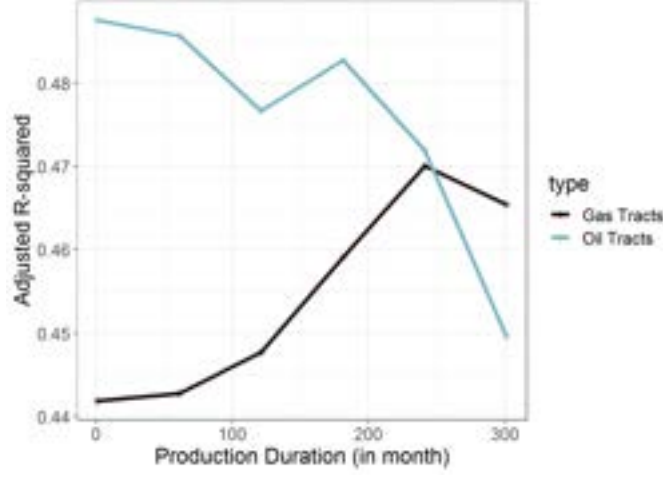
where $i \in \{o, g\}$, and $Q_t^{i,\tau}$ is the total amount of reserve type i produced by tract t after τ months of production, and X_t are tract characteristics. I allow the regression coefficients $\{\beta_0, \beta_{j,j=1,2,\dots,\underline{T}}, \alpha, \gamma\}$ to be dependent on the range of T . In my benchmark estimates, $\underline{T} = 2$ years, and I estimate regressions separately for $T < 5$ years, $T < 10$ years, ..., up to 30 years. Table 38 summarizes the number of completed and uncompleted tracts in my data, and figure 39 shows the model fit of (27).

Table 38: Summary of tract types

Tract type	N obs	Average Production Duration (in month)
Uncompleted oil tracts with sufficient data	770	335.4
Completed oil tracts	691	157.6
Uncompleted oil tracts without sufficient data	56	9
Uncompleted gas tracts with sufficient data	724	335.4
Completed gas tracts	2786	147.9
Uncompleted gas tracts without sufficient data	40	10.7

For tracts that were explored but not subsequently developed, their production quantities are assumed to be 0 ($Q^o = Q^g = 0$).

Figure 39: Oil and gas production model fit



Constructing $(\mu_{Gamma,oil}, \mu_{Gamma,gas}, \mu_{Bernoulli,oil}, \mu_{Bernoulli,gas})$ I decompose the prior F^Q into two priors F^{Q^o} and F^{Q^g} for oil and gas quantities, respectively. For each deposit type $i \in \{o, g\}$, the unconditional (with respect to any ex post outcomes of nearby tracts) prior $F_t^{Q^i}$ of a tract t is assumed to be a mixture distribution as follows:

$$\begin{cases} Q_t^i = 0 & \text{with probability } \frac{\exp(\beta_0^i + \nu_{0t}^i)}{1 + \exp(\beta_0^i + \nu_{0t}^i)} \\ Q_t^i \sim Gamma(\cdot, \cdot) & \text{with probability } \frac{1}{1 + \exp(\beta_0^i + \nu_{0t}^i)} \end{cases} \quad (28)$$

where β_0^i is a parameter to be estimated. ν_{0t}^i is assumed to be a Gaussian random field with a Matérn correlation function, where, for two tracts t and t' , the covariance between ν_{0t}^i and $\nu_{0t'}^i$ is

$$\sigma_0^{i2} (\gamma_0^i ||s_t - s_{t'}||) K_1(\gamma_0^i ||s_t - s_{t'}||), \quad (29)$$

where σ_0^i and γ_0^i are parameters to be estimated, $||s_t - s_{t'}||$ denotes the distance between the centroid of tract t and tract t' , and $K_1(\cdot)$ is the modified Bessel function of the second kind with a smoothness parameter of one (Krainski et al., 2018). Next, the Gamma distribution in Eq. (28) is assumed to have mean μ_t^i and standard deviation $\frac{(\mu_t^i)^2}{\phi^i}$, where

$$\log \mu_t^i = \beta_1^i + \nu_{1t}^i,$$

in which ϕ^i and β_1^i are parameters to be estimated, and ν_{1t}^i is also a Gaussian random field similarly defined as in (29) with parameters σ_1^i and γ_1^i . To estimate γ_0^i , I estimate the range value $\rho_0^i = \frac{\sqrt{8}}{\gamma_0^i}$. I follow a similar procedure for γ_1^i —i.e., we estimate $\rho_1^i = \frac{\sqrt{8}}{\gamma_1^i}$.

Thus, in summary, the parameters that need to be estimated are $\{\beta_0^i, \sigma_0^i, \rho_0^i, \phi_1^i, \beta_1^i, \sigma_1^i, \rho_1^i\}$ for $i \in \{o, g\}$. These parameters represent the spatial distribution of oil and gas deposits. I estimate these parameters using R-INLA (Rue et al., 2009). The estimated parameters of the priors are summarized in Table 40.

Table 40: Estimated parameters of the priors

	β_0^o	β_1^o	ϕ_1^o	ρ_1^o	σ_1^{2o}	ρ_0^o	σ_0^{2o}
Estimate	-7.112	14.750	0.388	2.203	1.955	1.492	2.944
SE	1.364	0.851	0.017	0.726	0.380	0.300	0.404

	β_0^g	β_1^g	ϕ_1^g	ρ_1^g	σ_1^{2g}	ρ_0^g	σ_0^{2g}
Estimate	-6.233	17.430	0.716	0.207	1.012	1.837	3.321
SE	1.207	0.073	0.019	0.041	0.077	0.304	0.428

The prior distribution of a particular tract t being auctioned in year τ is then the distribution in (28) conditional on ex post production and exploration outcomes of all tracts that were explored prior to year τ . This naturally implies that tracts that were auctioned in earlier years have less informative priors than those that were auctioned in later years. In addition, the same tract can have different priors based on the year of the auction.

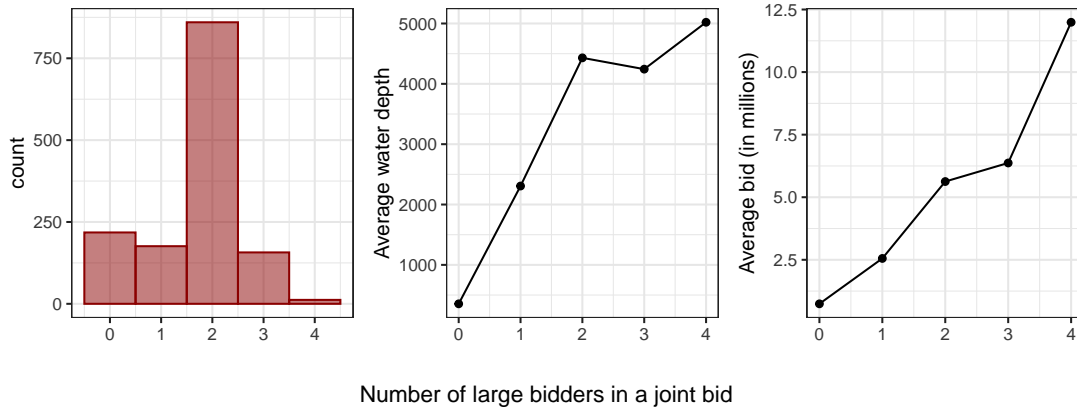
F Joint Bidding

In this section, I examine bidders' incentives to form joint bids. The analysis focuses on three explanations for joint bids proposed in the literature (see, for example, Hendricks and Porter (1992)): (1) alleviating capital constraints required for bidding, (2) reducing competition, and (3) sharing information among bidders. Notably, changes in the bid disclosure policy could influence the composition of joint bids if bidders form them to pool information.

To distinguish between bidders who are likely to possess private information on production values of tracts and bidders who primarily provide capital without directly engaging in drilling operations, I categorize bidders into two categories: large and small. Large bidders typically possess the expertise required for the exploration and development of offshore wells, allowing them to submit bids independently. In contrast, small bidders are less likely to make solo bids. A large bidder is defined based on their bidding activity from 2000 to 2019, considering both individual and joint bid contributions, exceeding the median bidder's

activity. This classification identifies 58 large bidders and 237 small bidders. Figure 41 illustrates the composition of joint bids. 28% of joint bids involve only one or no large bidders, and the majority of joint bids involve two large bidders. Joint bids among large bidders are more common in deeper tracts where bids are more expensive.

Figure 41: Composition and characteristics of joint bids



Due to the observation that some large bidders rarely compete against each other, I further organize them into clusters using a network approach, where each bidder represents a node and those participating in joint bids are linked. Using the algorithm in Blondel et al. (2008), I form 23 distinct clusters of bidders. In most auctions (65.34%), bidders do not compete against firms within their own clusters.

Table 42 summarizes the auction statistics for large bidders within and across clusters. The summary statistics suggest that (1) joint bids formed between large firms within the same cluster are likely intended to reduce competition rather than to pool information; (2) joint bids formed *across* clusters are more likely aimed at pooling capital; and (3) joint bids between large firms tend to target tracts that are ex ante less risky, which does not support the hypothesis that bidders form joint bids to share information.

Column (1) of Table 42 shows that within-cluster joint bids often occur on shallower, less expensive tracts that are ex ante gas-productive. For instance, the predicted probability of finding gas for these tracts is 29%, significantly higher than those for other tracts, where predictions range between 10% and 13%. This higher predicted probability of finding gas also implies a higher ex post probability of discovering gas. In comparison, when firms from the same cluster submit separate bids for the same tract (column (2)), the average winning bid is significantly higher (\$6.79M compared to \$3.19M). These tracts tend to be deeper, attract more bids, are more expensive, and have lower predicted probabilities of discovering

Table 42: Summary statistics of bids involving at least two large firms

	No within-cluster rival bidding, no across-cluster joint bids (1)	Within-cluster rival bidding, no across-cluster joint bids (2)	Across-cluster joint bids (3)
Number of auctions	1419	550	336
Average winning bid	3.19 (7.43)	6.79 (14.28)	12.99 (19.18)
Average losing bid	0.86 (2.59)	1.62 (3.93)	3.32 (5.58)
Average number of clusters	1.81 (0.65)	1.85 (1.06)	3.19 (0.97)
Average number of bids	2.34 (0.7)	3.2 (1.49)	3.75 (2.12)
Average water depth	2966.21 (2921.97)	4416.64 (2280.69)	4867.47 (2532.52)
Bernoulli prior (oil)	0.12	0.07	0.1
Bernoulli prior (gas)	0.29	0.1	0.13
Productive in oil	0.04	0.03	0.04
Productive in gas	0.12	0.04	0.03

Note: Numbers in parentheses represent the sample standard deviation. Bid statistics are reported in millions of dollars.

gas and oil than tracts in Column (1).

Column (3) of Table 42 also indicates that large bidders form joint bids *across* clusters for highly competitive and expensive tracts that are likely to contain more oil and gas than tracts where bidders from the same cluster submit separate bids. In addition, the average bid for these tracts is more than double the winning bids for other tracts in the sample. This result is consistent with the hypothesis that firms form joint bids to pool the capital required for bidding.

Table 43 analyzes the propensity of bidders to submit joint bids, controlling for firm and tract location fixed effects. In columns (1) and (2), I examine whether firms are more likely to submit joint bids for tracts that turn out to be more productive, considering all bids and those submitted by large bidders, respectively. In column (3), I investigate whether a positive correlation exists between the decision of large firms to form joint bids with other large firms in different clusters and the productivity of a tract. The results indicate that joint bids are more likely to be formed for tracts that are ex ante more productive. For instance, a

Table 43: Determinant of joint bids

	<i>Dependent variable:</i>		
	Joint bids		Joint bids across clusters
	(1)	(2)	(3)
Bid (in log)	0.053 (0.004)	0.053 (0.004)	0.024 (0.004)
Number of bids	0.002 (0.003)	0.004 (0.003)	0.004 (0.002)
Tract's water depth	−0.008 (0.010)	0.003 (0.012)	0.002 (0.009)
Productive in oil (ex-post)	−0.001 (0.028)	−0.013 (0.032)	0.005 (0.026)
Productive in gas (ex-post)	0.010 (0.019)	−0.013 (0.026)	−0.059 (0.015)
Productive in oil (ex-post) x winning	−0.014 (0.045)	0.006 (0.053)	0.037 (0.048)
Productive in gas (ex-post) x winning	0.012 (0.028)	−0.030 (0.039)	0.030 (0.026)
Bernoulli prior (gas)	−0.007 (0.056)	−0.050 (0.068)	0.101 (0.049)
Gamma prior (gas)	0.0001 (0.012)	−0.018 (0.017)	−0.015 (0.010)
Bernoulli prior (oil)	0.159 (0.074)	0.160 (0.086)	0.015 (0.070)
Gamma prior (oil)	0.033 (0.011)	0.019 (0.013)	0.029 (0.009)
Constant	−1.154 (0.266)	−0.676 (0.343)	−0.669 (0.209)
Firm FEs	Yes	Yes	Yes
Area FEs	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes
Sample	All	Large bidders	Large bidders
Observations	8,549	6,673	6,673
Adjusted R ²	0.352	0.330	0.159

Note: In the first column, the sample includes all bids from auctions held between 2000 and 2019 that involved at least two bidders. The second and third columns focus exclusively on bids involving at least one large bidder. The dependent variable in columns (1) and (2) is a binary indicator of whether a bid is a joint bid. In column (3), the dependent variable is a binary indicator of whether a joint bid was formed between large firms across different clusters. Standard errors, adjusted for heteroskedasticity, are reported in parentheses.

1% increase in the prior probability of a tract containing oil raises the likelihood of forming a joint bid—either between any two firms or between two large firms—by 0.16 percentage points. Additionally, firms are more inclined to submit joint bids when planning to place large bids. Specifically, a 1% increase in bid value corresponds to a 0.05 percentage point increase in the probability of submitting a joint bid. These findings align with the theory that firms form joint ventures to pool capital for bidding on competitive tracts, consistent with the results of Hendricks and Porter (1992). However, the post-bid outcomes—such as whether a tract was explored or produced oil or gas—do not exhibit a significant positive correlation with joint bid status across all regression models. This suggests that joint bidders likely do not possess an informational advantage over solo bidders.

F.1 A model of joint bidding with communication friction

In this section, I present a stylized model to examine firms' incentives to form joint bids. In this framework, bidders can form joint bids but incur a cost, as doing so requires them to reveal private information about the tract's production potential. This modeling choice is motivated by the fact that firms in the OCS setting are often reluctant to share proprietary information related to their technology, which determines the profitability of a tract. I first analyze the equilibrium joint bid formation under the full disclosure (FD) policy, followed by the non-disclosure (ND) policy.

Consider a first-price sealed-bid auction with two bidders, where each bidder i (for $i = 1, 2$) receives a private signal S_i , drawn uniformly from the interval $[0, 1]$. After the auction, the winner has the option to incur a cost C to extract the value of the tract, which is given by:

$$V = \frac{S_1 + S_2}{2}.$$

If the winning bidder chooses not to pay this cost, the value of the tract is 0. The cost C is drawn from a uniform distribution over the interval $[0, 1]$, i.e., $C \sim U[0, 1]$.

Bidders also have the option to participate in a joint bidding venture. To bid jointly, each bidder i must disclose a message M_i , which is defined as $M_i = \sigma_i \epsilon_i + (1 - \sigma_i) S_i$, where ϵ_i is randomly drawn from $U[0, 1]$. Here, σ_i , which can be either 0 or 1, denotes the precision of M_i . When $\sigma_i = 0$, bidder i fully shares their private information with the joint bid partner, while $\sigma_i = 1$ indicates that the message M_i contains no meaningful information about the tract. The choice of σ_i is made privately before the realization of the private signal S_i and before forming the joint bid. I assume that choosing $\sigma_i = 1$ incurs a cost of $\underline{\kappa}$, while choosing

$\sigma_i = 0$ incurs a higher cost, $\bar{\kappa}$, where $0 \leq \underline{\kappa} < \bar{\kappa}$. In a joint venture, the combined bid and post-auction decisions are determined by the aggregate information provided by both parties, which is equal to $S_J = \frac{1}{2}(M_1 + M_2)$.

In this example, $\underline{\kappa}$ represents the cost of forming a joint venture, which could arise if the government deems the auction non-competitive and rejects the joint bid. In a more general model with multiple bidders, this cost can also capture (in a simplified way) the equilibrium effect of increased equilibrium bids when all bidders decide to form joint bids.²⁹

Throughout this exercise, I focus on symmetric Bayes-Nash equilibria. In the following analysis, I examine the equilibrium decisions regarding whether to bid jointly and the precision of M_i under both the FD policy and the ND policy for losing bids.

Full-disclosure policy

Under an FD policy of losing bids, let $\beta_{solo}^{FD}(\cdot)$ denote the equilibrium bidding strategy. Firm i 's profit from bidding b is as follows

$$\mathbb{E} \left(\int^{\frac{1}{2}(s+s')} \left[\frac{1}{2}(s + s') - c \right] dc - b | s' \leq \beta_{solo}^{FD^{-1}}(b) \right) \Pr \left(s' \leq \beta_{solo}^{FD^{-1}}(b) \right)$$

Applying Proposition 1 in the main text, conditional on $S_i = s$, bidder i 's bidding strategy is given by

$$\beta_{solo}^{FD}(s) = \frac{1}{6}s^2$$

The expected profit from solo bidding is given by

$$\pi_{solo}^{FD}(s) = \frac{3}{24}s^3$$

Now, let's consider the incentive to form a joint bid. When both firms form a joint bid, this becomes an auction with only one bidder. Therefore, the joint bidder will bid $b = 0$ and obtain the tract.

²⁹For example, Krishna and Morgan (1997) studies equilibrium bidding strategies in a common-value auction (without ex post actions), assuming the latent value is the average of all firms' signals. The key result of this paper is that joint bidding mitigates the winner's curse and leads to higher equilibrium bids with fewer bidders, under the assumption that the signal distribution is log-concave. DeBrock and Smith (1983) reach a similar conclusion using numerical simulations. However, these papers do not examine whether forming a joint bid is optimal.

Conditional on $\sigma_j = 0$, the expected profit of firm i from joint bidding is given by

$$\mathbb{E} \left(\frac{1}{2} \int^{\frac{1}{2}(M_i+S_j)} \left[\frac{1}{2}(S_i + S_j) - C \right] dC \right) - \kappa(\sigma_i) = \begin{cases} \frac{s}{8} - \underline{\kappa} & \text{if } \sigma_i = 1 \\ \frac{s^2}{16} + \frac{s}{16} + \frac{1}{48} - \bar{\kappa} & \text{if } \sigma_i = 0 \end{cases} \quad (30)$$

Conditional on $\sigma_j = 1$, the expected profit of firm i from joint bidding is given by

$$\mathbb{E} \left(\frac{1}{2} \int^{\frac{1}{2}(M_i+M_j)} \left[\frac{1}{2}(S_i + 0.5) - C \right] dC \right) - \kappa(\sigma_i) = \begin{cases} \frac{s}{8} - \frac{1}{96} - \underline{\kappa} & \text{if } \sigma_i = 1 \\ \frac{s(s+1)}{16} + \frac{1}{96} - \bar{\kappa} & \text{if } \sigma_i = 0 \end{cases} \quad (31)$$

Proposition 4 *Under FD:*

- Bidders do not form a joint bid in equilibrium if $\underline{\kappa} \geq \frac{1}{96}$ and $\bar{\kappa} \geq \frac{1}{24}$.
- If $\underline{\kappa} < \frac{1}{96}$ or $\bar{\kappa} < \frac{1}{24}$, bidders form a joint bid and choose not to reveal their private signals if $\frac{1}{96} \leq \bar{\kappa} - \underline{\kappa}$.
- If $\underline{\kappa} < \frac{1}{96}$ or $\bar{\kappa} < \frac{1}{24}$, bidders form a joint bid and choose to reveal their private signals at cost of $\bar{\kappa}$ if $\bar{\kappa} - \underline{\kappa} < \frac{1}{96}$.

Proof:

From (30), it can be shown that firm i will choose $\sigma_i = 0$ (1) if $\bar{\kappa} - \underline{\kappa} \leq (>) \frac{1}{96}$. Similarly, the same threshold applies to (31). Therefore, the equilibrium choices of (σ_i, σ_j) are as follows:

- $(\sigma_i = 0, \sigma_j = 0)$ if $\bar{\kappa} - \underline{\kappa} \leq \frac{1}{96}$, yielding an expected equilibrium profit for each firm of $\pi_{joint} = \frac{7}{96} - \bar{\kappa}$.
- $(\sigma_i = 1, \sigma_j = 1)$ if $\bar{\kappa} - \underline{\kappa} > \frac{1}{96}$, yielding an expected equilibrium profit for each firm of $\pi_{joint} = \frac{5}{96} - \underline{\kappa}$.

The ex-ante expected profit from solo bidding under FD is $\mathbb{E}(\pi_{solo}^{FD}(s)) = \frac{3}{96}$. Therefore, the statements of Proposition 4 are obtained. ■

Non-disclosure policy

Under an ND policy, let $\beta_{solo}^{ND}(\cdot)$ denote the equilibrium bidding strategy. Firm i 's profit from bidding b is as follows

$$\left(\int^{\mathbb{E}\left(\frac{s+s'}{2} | s' \leq \beta_{solo}^{ND-1}(b)\right)} \left[\mathbb{E} \left(\frac{s+s'}{2} | s' \leq \beta_{solo}^{ND-1}(b) \right) - c \right] dc - b \right) \Pr \left(S_j \leq \beta_{solo}^{ND-1}(b) \right)$$


Applying Proposition 2 in the main text, conditional on $S_i = s$, bidder i 's bidding strategy is given by

$$\beta_{solo}^{ND}(s) = \frac{9}{64}s^2$$

The expected profit from solo bidding is given by

$$\pi_{solo}^{ND}(s) = \frac{9}{64}s^3$$

Figure 44: Illustrations of the optimal bidding and messaging strategies



Note: The gray region represents the values of $(\bar{\kappa}, \underline{\kappa})$ for which bidders will bid solo in equilibrium. In the blue region, bidders bid jointly but disclose only uninformative messages. In the red region, bidders bid jointly and fully reveal their private signals.

Proposition 5 *Under ND:*

- Bidders do not form a joint bid in equilibrium if $\underline{\kappa} \geq \frac{5}{768}$ and $\bar{\kappa} \geq \frac{29}{768}$.
- If $\underline{\kappa} < \frac{5}{768}$ or $\bar{\kappa} < \frac{29}{768}$, bidders form a joint bid and choose not to reveal their private signals if $\frac{1}{96} \leq \bar{\kappa} - \underline{\kappa}$.
- If $\underline{\kappa} < \frac{5}{768}$ or $\bar{\kappa} < \frac{29}{768}$, bidders form a joint bid and choose to reveal their private signals at a cost of $\bar{\kappa}$ if $\bar{\kappa} - \underline{\kappa} < \frac{1}{96}$.

Proof:

The proof is similar to that of Proposition 4 and is therefore omitted. ■

Figure 44 illustrates Propositions 4 and 5. Here, bidders are *less* likely to form a joint bid after a non-disclosure policy is implemented. This is because the profit under solo bidding is higher under ND than under FD. Despite being able to obtain the tract at zero bid as a joint venture, forming a joint bid is costly to the bidders for two reasons. First, when firms choose not to reveal their private signals, the post-auction decision of whether to incur C is suboptimal. Second, the cost of revealing more information is entirely borne by the bidder, whereas some of the gain from this additional information is expropriated by the other bidder.