Investment and Taxation: the Case of Oil and Gas in the Permian Basin

Andrey Ordin

Duke University, Department of Economics

October 31, 2019

ABSTRACT. This paper evaluates the sensitivity of investment and production for the upstream oil industry in the Permian Basin to taxes. I set up a model of auctions for land leases followed by well drilling by the lease holder. Leveraging the strategic decisions of agents about the timing of drilling, the model allows me to evaluate the distortionary effects from revenue taxes in the industry even in the absence of variation in these taxes. I estimate the model with data from the state of New Mexico and use it to show how alternative policies can affect oil production and state revenues. The results show significant distortions in drilling activity and moderate scope for improving the state's revenues through alternative tax policies.

Latest version of the paper here

CONTACT. andrey.ordin@duke.edu.

I am deeply grateful for the valuable comments and support in research from James W. Roberts, Allan Collard-Wexler, Daniel Yi Xu, Christopher D. Timmins, Juan Carlos Suárez Serrato, and Vivek Bhattacharya. I also thank the participants of the Duke IO/Labor Lunch Group for the useful feedback on the project. All errors are my own.

1. Introduction

In September 2019 the US Government Accountability Office published a report saying that the government missed out on $18bln from firms producing oil in the Outer Continental Shelf. Normally, these firms would pay a fraction of their revenues, in the form of royalty payments, to the owner of the lands on which they operate; however, between 1996 and 2000 the federal government sold a number of offshore oil leases without such a revenue sharing clause.[1] Why was royalty lifted in the first place? As a form of revenue tax, it distorts the investment decisions of the firms; thus, the goal of the government was to incentivize drilling in the region of interest. Around the world, however, royalties are a popular instrument for collecting revenues, in some countries reaching the overwhelming level of 90% (Goldsworthy and Zakharova (2010)). Both the international experience and the US example suggest that, in general, lease contract terms imply a trade-off between the level of economic activity and the payoff of the seller. How should these contracts be designed? What policies, perhaps akin to royalty relief, should the government pursue in this industry, and how can they affect firm investment decisions and government revenues? The aim of my paper is to answer these questions.

To understand how tax policy drives the incentives to drill for oil, I look at the Permian Basin region of south-eastern New Mexico. The Permian Basin is one of the most productive and oil-rich areas in the world, its output outpacing that of entire countries such as Libya or Oman. With firms eager to extract the mineral resources, the New Mexico State Land Office, responsible for leasing state lands to firms, reports[2] that in 2019 alone its activities will generate estimated $1,500 per beneficiary taxpayer. Yet, despite the importance of the industry, today the state still uses tax policies and lease terms that it relied on 25 years ago. In this setting, I analyze how firms would respond to alternative policies by developing a structural model of the leasing process followed by a dynamic drilling decision. I use the model to quantify the distortions from the revenue taxes already in place. Furthermore, the model allows me to study various taxation schemes, from simple revenue-side and cost-side taxes and subsidies to sliding scale royalties similar to the relief program described above, and to scenarios where firms themselves compete for leases by bidding taxes.

[1] The royalty was supposed to be relieved only when oil prices stayed below $35, but some of the lease contracts omitted this contingency, leading to losses that continue to accumulate over the production period. The omission led to public outcry and was discussed on congressional hearings (see Haile, Hendricks, and Porter (2010)).

[2] NM State Land Office link, accessed October 2019.

In New Mexico, as well as all around the US, public lands are leased by the government agencies primarily via first-price cash auctions. The highest bidding firm receives a right, though not an obligation, to drill for oil within a finite time period. In absence of drilling, eventually, the lease expires. Traditionally, firms with producing wells have been exposed to a variety of taxes, including in particular royalties and severance imposed on the revenue streams from the wells. These revenue taxes historically have seen very little variation: for instance, royalties for federal leases have largely remained at the 12.5% level since 1916.[3] At the same time, for top producing states, oil and gas revenues collected as taxes and cash bids for leases can be of major importance. In North Dakota, for example, these revenues constitute about $2bln annually, or almost 50% of the general revenue fund. Indirect effects in the form of job creation and various local spillovers (Hausman and Kellogg (2015), Allcott and Keniston (2017)) further emphasize the value of the industry for the states. It is therefore not surprising that proposals are being made to implement alternative tax schemes for public land leases, ranging from fairly simple readjustments (Rabe and Hampton (2015)) to more sophisticated tools such as a sliding scale royalty (see, e.g., Bureau of Land Management (2015)).

Classical models for auctions could possibly be used to analyze the bidding response to the rare shifts in the government policy (Hendricks and Porter (2014)), but such models are not set up to describe what happens to the investment in the industry. For standard approaches to policy evaluation, the lack of changes in the said policy is a significant impediment. Here, I develop methods which make the analysis possible. There are two major challenges I need to overcome in this setting. First, variation over time in revenue tax rates imposed by the states is typically very low, if present at all. Second, to understand the investment incentives of the firms it is necessary to evaluate the profitability of wells, both the ones that exist and the ones that could potentially have been drilled on the leased lands. Productivity of wells can vary dramatically from one oil field to another, and, similarly, the effective costs of drilling show significant geographic differences (Energy Information Administration (2016)). Therefore, potentially unobservable regional features characterizing productivity and the investment costs of the firms are important for policy design. The structural model I set up helps me to address both of these problems.

To deal with the lack of tax variation, I use the price of oil, which is a major driver of investment

[3] One exception is the windfall profit tax that was in effect between 1980 and 1986 in response to the rise in oil prices — see Rao (2018) for details.

incentives. In my model, although the prices themselves are exogenous, the decision *when* to drill a well is made by the firm holding the lease. The choice of this timing is cast as an optimal stopping problem in an environment with stochastically evolving oil prices. The threshold price that induces the decision to invest depends on per-unit costs of drilling and the time remaining till expiration of the lease. The spud dates of the wells observed in the data allow me to determine per-unit costs. Using expected lifetime output levels for each of the wells in the sample, I show how state revenues from taxes change as the tax rates go up. Under the optimal stopping assumptions, rising taxes delay the drilling of each of the wells, eventually driving them out of existence. On the other hand, the state gets to collect more from the wells that do get drilled. For my data, this trade-off implies that the tax revenue can be improved only slightly over the baseline, and that significant production losses start to accumulate as the tax rate exceeds the 30% threshold.

To account for other channels of state revenues and to consider more sophisticated policies, I augment this model with a stage where firms compete for lands in a first-price auction. Each firm receives a noisy signal about the quantity of oil contained underground; the quantity is common to all the agents. Signals inform the firms' expectations about the option value of holding the lease, which, in turn, is founded on the optimal stopping problem. Based on the signals, auction participants submit their bids and the highest bidder wins the rights to the lease. This two-stage model has several useful features. First, it allows me to determine the primitives defining perceived profitability of leased lands for the firms, namely the associated quantities of oil and costs of extraction. Second, it helps me to investigate the effects of policies interacting with some of these primitives, such as revenue taxes or cost subsidies. Third, as discussed below, I can use it to study a different way to set up the lease terms: let the firms bid in tax rates rather than cash, similarly to how they bargain over royalty rates with private land owners.

I estimate this model using bidding and drilling data from the Permian Basin region in New Mexico. I begin my analysis by determining what happens if the distortionary revenue taxes are lifted. I find that while cash bids go up, overall state revenues decrease. On the other hand, economic activity improves: the frequency of drilling goes up by about 60%, though the total industry profits rise only by 8%. In other words, with removal of the distortions, a host of marginally profitable wells get drilled and the government share in profits from extraction goes down.

Looking at simple revenue-side and cost-side policies, I find that the state cannot significantly improve its revenues using simple adjustments to its programs. Indeed, the best-performing tax at the rate of 30% boosts revenues by less than 3% while bringing down expected drilling rates. This result is consistent with the exercise described above, where the data for existing wells was used in the optimal stopping setting. Another revenue-side policy I consider is inspired in part by the federal royalty relief program. Just like in that program, I let the revenue tax depend on prices, and I look for the dependence that would maximize the payoff of the state. I find that such a sliding scale can perform very well: it is possible to increase the revenues of the government by about 10% at the cost of slightly reducing the drilling incentives of the firms. Extrapolating for all government collections from oil and gas in New Mexico, this improvement would correspond to about $200mln of extra revenues each year.

While royalties associated with leases for public lands are fixed by the seller and largely do not vary, royalties for private land leases are set through negotiations between the landlords and the extraction firms.[4] Having the auction stage in the model allows me to predict outcomes in a counterfactual scenario where, similarly, firms submit revenue share bids rather than cash. As the economic literature suggests, how such auctions compare to simple cash in terms of seller revenues is an empirical question (Che and Kim (2010)). My results indicate that in the case of New Mexico tax bidding would lead to a dramatic drop in economic activity. Without incurring particularly significant losses in the absence of oil production, the firms end up submitting high revenue fractions as their bids, gambling on high future oil prices and undermining their own incentives to drill. The total industry profits and the revenues of the state end up getting slashed in half.

It is well understood in the economic literature that in environments with ex-post uncertainty profit sharing agreements perform very well for maximizing the revenues of the seller without eroding the investment decisions of the firms (Cong (2019)). In line with this observation, my model predicts that the state can soak up almost all of the industry's profits by implementing such agreements, diverting the firms from bidding for leases only with tax rates above 90%. My analysis also shows that a 40% profit tax[5] is equivalent in state revenue to the baseline, but superior to it in terms of the drilling activity.

The rest of the paper is organized as follows. After discussing how my paper fits into the

[4] In the state of Louisiana, public land leases are sold via auctions with two-dimensional bids that have both upfront cash and revenue share components — see Kong, Perrigne, and Vuong (2019) for details.

[5] For comparison, in Alaska the local government currently uses a 35% profit tax.

literature, I proceed to Section 2 where I describe the empirical setting and the data. Section 3 introduces the model for optimal investment timing and shows how it can be used to evaluate effects from increases in revenue taxes. Section 4 sets up the preliminary first-price sealed bid auction stage and shows how the basic predictions of the model are borne out in the data. The estimation strategy is described in Section 5. Finally, Section 6 considers the counterfactual scenarios of alternative tax policies and Section 7 concludes.

Related Literature

There are several major strands of research this paper is related to. First, much work has been done on the subject of auctions for oil and gas leases. Here, papers analyse differences across bidders (Hendricks and Porter (1988)), the capacity of theoretical models to describe the bidding decisions of the agents observed in practice (Hendricks, Pinkse, and Porter (2003)), and the matters of auction design (Hendricks, Porter, and Tan (1993)). More recent works study the synergy effects from winning multiple auctions (Kong (2017b)). Importantly, however, my model involves both competition in auctions and strategic activity that follows. Recently, a series of empirical works with similar two-stage frameworks (Bajari and Lewis (2011), Bajari, Houghton, and Tadelis (2014), An and Tang (2019)) have found that in these settings the design of the underlying contracts can be highly important for the revenues of the seller and the efficiency of the outcomes.

Second, a range of works in public economics analyze the links between investment and taxation. A line of research stems form the seminal paper of Hall and Jorgenson (1967) where businesses are demonstrated to actively readjust the volume and timing of their investment in response to various tax incentives. Works on bonus depreciation in particular (e.g., House and Shapiro (2008)) study policies similar to the intangible cost deduction rule for oil and gas wells that motivates my cost-side analysis. Other related research looks at the incidence of tax incentives (Goolsbee (1998)), heterogeneity in response of the firms (Zwick and Mahon (2017)), or the overall effect of corporate taxes on investment and growth (Lee and Gordon (2005), Djankov et al. (2010)). Unlike these works, I leverage the power of structural modeling to infer investment incentives of the firms across different projects in absence of tax variation.

Third, the oil and gas industry itself has received a lot of interest from researchers. Here, papers study how upstream oil firms respond to changes in their incentives to produce. A key result from

Anderson, Kellogg, and Salant (2018) is that firms do not readjust their well production curves over time. Other works measure the response of the industry on the extensive margin: Newell, Prest, and Vissing (2016) and Metcalf (2018) evaluate the elasticity of drilling probability in oil prices, and a series of papers determine how output and drilling change in response to shifts in severance tax rates of the states (e.g., Reimer, Guettabia, and Tanaka (2017), Brown, Maniloff, and Manning (2018)). While I also use variation in prices, I rely on a structural model of investment timing to analyze how it affects the drilling decisions. This approach allows me to work in a setting that has virtually no variation in revenue taxes. Additionally, my model allows me to evaluate the potential impact from counterfactual taxation policies of varying complexity. While I determine investment cost structure that would rationalize the strategies of the firms observed in the data, Asker, Collard-Wexler, and Loecker (2018) demonstrate how cost data, when available, can be used to understand inefficiencies in oil production decisions. The authors evaluate the discrepancy, or the wedge, between the socially optimal and the observed allocation of production around the world and quantify how various factors contribute to it. One of such factors is revenue taxes that I study here.

Finally, a recent series of empirical papers analyze the links between economic activity and design of contingent payment contracts in the oil and gas setting. Contingency here refers to the dependence of the landlord's payoffs on the revenue stream of the buyer. Herrnstadt, Kellogg, and Lewis (2019) study private leases for gas extraction in the Haynesville shale of Louisiana. The authors build a principal-agent model to discuss the implications of alternative contract design for private landowners, finding that lease duration and the associated royalty can be used as imperfect tools to extract revenues from the firm. Kong, Perrigne, and Vuong (2019) look at oil and gas lease auctions in Louisiana where, akin to private leasing, two-dimensional bids are used to assign the lands. In their unusual setting for empirical auctions, participating firms offer both upfront cash and a fraction of future revenues as their bids. The authors find that the standard fixed-royalty lease terms would lead to much better outcomes for the industry as well as for the state. Bhattacharya, Ordin, and Roberts (2019) consider contingent payment auctions where participants compete for the leases using security bids rather than plain cash bids. The paper evaluates the potential impacts from the use of such alternative allocation mechanisms. By contrast, in this paper I study the state policies aimed at the producing wells and, as part of my analysis, investigate what the relationship between well delays and their expected productivity implies for potential policy outcomes.

The empirical works mentioned above together with theory results for settings similar to mine (e.g. Board (2007), Cong (2019)) suggest that the current practices of New Mexico in public leasing lead to investment and production distortions. I evaluate the magnitude of these distortions and determine which alternative taxation policies could improve the revenue of the state without hurting the economic activity in the industry.

2. Empirical Setting and Data

In this section, I describe the setting and the data used in the analysis.

I analyze auctions held by the New Mexico State Land Office (NMSLO), which manages 9 million acres of surface and 13 million acres of subsurface estate for the beneficiaries of the state land trust such as schools, universities, and hospitals. The NMSLO sells leases that grant the right to drill for oil and gas. The proceeds are put in the Land Grant Permanent Fund, which is invested for the beneficiaries. The Land Office generates hundreds of millions of dollars each year while saving the average beneficiary $1,500 in annual taxes.

The sales are held each month on a fixed schedule, with the goal of leasing roughly 10,000 acres per month. The tracts can be nominated by the firms, but the exact tracts up for sale are published by the office only a short time (currently one month) in advance. Though recently the state switched to holding the auctions online, in the time period relevant for the sample the participants had to send in their bids in sealed envelopes. They also had to pay for and fill out the participation documents, though this fee is negligibly low (currently $30). The firm producing oil and gas must be registered with the taxation and revenue department of New Mexico to participate in public leasing activities. This being the only restriction on participation, most of the firms competing for state leases are relatively small. Aside from bidding, they either find an operator or operate the wells, though they rent the rig and hire the crews to do the drilling itself.

The leases are sold via cash auctions with a fixed royalty component. First, all bidders submit cash bids, also known as bonuses. The bidder who submits the highest bonus wins the auction and pays the seller this bid immediately following the auction. The second component involves royalty payments. The royalty rate is set ahead of time for each lease by NMSLO, with the standard rates being 1/8, 3/16, and 1/6. The choice of the rate is loosely based on the geographic location of the lands. The state additionally levies various production taxes on all oil and gas wells with the total tax rate on the order of 7%. The auction winner receives a right, but not an obligation, to drill a well on the land within a five year period. If the well is drilled, the land is "held by production:" for as long as oil is extracted, the lease can be held by the winner beyond the original term. Finally, the auctions have a publicly known royalty rate of roughly $15.625 per acre.[6] This reserve rate is rarely binding in auctions and for a typical lease corresponds to the value of $5,000, over five times lower than the average bid in the sample. In what follows I ignore this rate and assume it to be non-binding, similar to how it is done in Kong (2017b), another work studying the NMSLO auctions.

In addition to revenue-side distortions, the upstream oil and gas industry has historically been exposed to a set of various deductions rules, some of which act in effect as a form of cost subsidy. One of the more impactful among these rules is the intangible cost deduction that has been made available to the firms operating around the country since 1916. For an oil well, drilling costs can be broken into two components: tangibles, i.e., recoverable costs such as drilling equipment, and intangibles, which cover everything else: labor, chemicals for injection, infrastructure for the area around the well. The latter constitute somewhere around 60%-80% of the total costs. Intangible costs are typically deductible from the tax base; however, while for many industries the deduction had to be distributed over 7 years, for the oil industry it has always been an option to deduct the entirety of the costs in the year when the investment is made. The Joint Committee on Taxation estimates that the federal government foregoes revenues on the order of $2bln each year because of this rule. Motivated by this example, as one of the applications for my model I consider the potential effects of similar cost-side policies that the state could pursue.

The dataset is put together from several sources.[7] From the NMSLO, I know the exact locations of the leased lands, the bids submitted at the auctions, and the names of the bidders.[8] The NMSLO uses both open outcry and sealed-bid auctions, though for estimation I focus on the sealed-bid auctions only. Based on the full sample, I construct a measure of potential bidders N for each auction following the approach traditional in the literature (e.g., Athey, Levin, and Seira (2011)).

[6] In the recent years, the minimum bid requirements have changed and may now depend on geographic location of the lease.

[7] This is the dataset used in Bhattacharya, Ordin, and Roberts (2019).

[8] For open outcry auctions, I only observe the winning bid and the name of the winner.

	Nobs	Mean	StDev	25%	50%	75%
			Auction features			
Oil price at auction P_0, \$/bbl	786	48.6	27.5	26.3	38.0	69.4
Acreage, acres	786	275	137	160	320	320
Profit tax, fraction	786	0.431	0.023	0.408	0.412	0.456
Revenue tax, fraction	786	0.232	0.012	0.237	0.237	0.237
			Bidding			
Winner's bid, \$ 1'000	786	65.9	72.0	18.2	40.2	87.0
Runner-up bid, \$ 1'000	506	42.0	48.1	12.0	26.0	52.7
Mean bid, \$ 1'000	786	41.5	37.6	15.2	29.9	58.1
$\text{Bid}_1/\text{Bid}_2$	506	3.37	4.86	1.3	1.86	3.35
Bid_1/P_0	786	1410	1173	469	1043	2127
Num bidders	786	2.3	1.4	1	2	3
Num potential bidders	786	4.0	2.3	2	3	5
			Post-auction outcomes			
$\mathbb{1}$(drilled)	786	.099	.3	0	0	0
Delay, days	78	1294	611	829	1556	1812
Q(oil), 1'000 bbl	54	75.4	96.8	6.3	30.4	135.1

Table 1: Sample summary, 1994-2012. All dollar values are normalized to their 2009 levels.

To do that, for every lease I count the number of unique bidders[9] for leases corresponding to adjacent lands sold via recently held auctions — see Appendix A for details. With Drillinginfo data on producing wells, I match the wells to leases based on their geographic boundary descriptions. Although Drillinginfo also provides data on the output of wells, due to partially missing observations I only use it to define oil-rich wells and to check model fit, but not otherwise. The oil price data used is monthly West Texas Intermediate crude oil spot prices.[10] On the state level, revenue taxes include royalties and production taxes[11], as well as a corporate income tax which is a form of profit tax. On the federal level, corporate income taxes, which in the model are viewed as a form of profit tax, apply. Thus, all the distortionary taxes are collected at the state level.

To homogenize the sample, I focus on the lands which were considered for oil rather than gas extraction. To do this, I first reduce the data to leases in the counties Lea and Eddy located in the Permian Basin region of the state. This, however, is not enough, since some of the lands even in these counties have been historically producing gas. Thus, I only keep the leases that are surrounded

[9]One bidder, Yates Petroleum, participates in these auctions so often that I include it as a potential bidder everywhere. The same approach was taken in Kong (2017b) and Bhattacharya, Ordin, and Roberts (2019).

[10]I compare the spot prices with the futures in Appendix B and find that the differences between the two price processes are fairly minor. In the main part of the paper, I use the spot prices.

[11]Production taxes include severance, emergency school fund, and ad valorem taxes. In New Mexico, royalties are deductible from production tax liability.

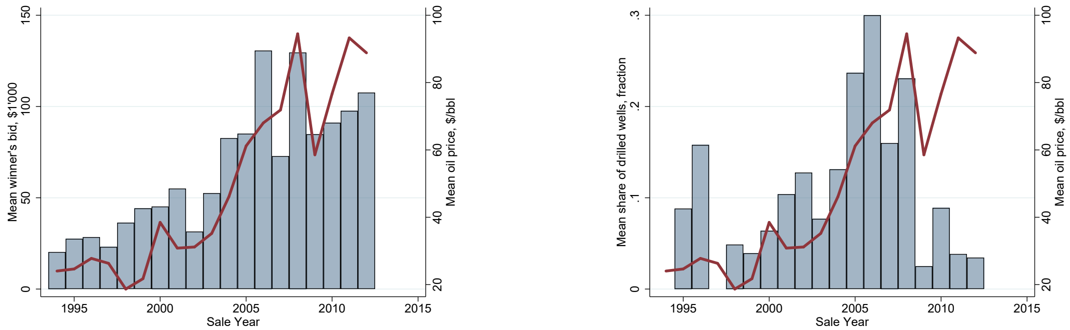

Figure 1: Mean annual winning bids (left) and fraction of drilled leases (right). Strategic decisions of the firms respond to oil prices.

by predominantly oil-rich wells — see Appendix A for details.

Summary statistics for the resulting sample are provided in Table 1. The data spans the years 1994 through 2012; over this time segment, many variables show significant variation. The third quartile of prices, for instance, is almost three times larger than the first quartile: \$69.4 and \$26.3 respectively. This variation is useful since price is the key variable representing the incentives to drill wells. Tax levels, on the other hand, are notably stable throughout the sample. Further, the summary shows large spreads in bids. The winning bid tends to be 50% larger on average and across the different quartiles than the runner-up's bid. This suggests significant disparity in how valuable the leases are to different bidders. Statistics for the ratio of winner's bid to oil price at auction also indicate that prices are a major source of differences in bids across auctions, though it leaves plenty of variation in the winner's bid unexplained. Participation in the auctions themselves tends to be rather low: more than half the sales saw not more than two competitors, and about 1/3 of all auctions involved a single bidder. On the post-auction side, the most striking feature is how infrequent drilling is: only 10% of all leases end up with wells. Delays for those wells are very long: almost four years on average. Such long delays can naturally arise as a feature of the solution to the optimal stopping problem discussed in Section 3.2.

The evolution of drilling decisions over time is presented in Figure 1. The two graphs demonstrate that bidders recognize the appreciation of land value with oil prices and that the drilling decisions of the lease holders are related to the movement of prices as well. These relationships are studied in more detail in the basic reduced-form analysis presented in Section 4.2. The structural model develops further insights into strategic response of agents to changes in price.

3. Optimal Investment Timing

In this section, I introduce a model describing the choice of well drilling timing and apply it to investigate the effects from raising taxes on state revenues.

3.1. An Optimal Stopping Model for Drilling Decisions

Consider an agent who holds a lease containing q barrels of oil. Over the time horizon of T years the agent has the opportunity to drill a well, extract the oil, give up a fraction ϕ of revenues to the state and keep the rest. At time t the costs of drilling are c_t; furthermore, the costs may depend both on the quantity of oil q and, contemporaneously,[12] on oil prices P_t. For simplicity, I assume that the decision to drill, production, and revenue transfers all occur simultaneously and instantly.[13] Thus, if the well is drilled at time τ, the contemporaneous payoff to the winner is

$$(1-\phi) \cdot P_\tau \cdot q - c_\tau. \tag{1}$$

The key insight here is that τ is endogenous: it is chosen by the lease holder based on expectations about the future behavior of oil prices and the profitability of the well.[14] Through this choice, the agent solves the maximization problem which gives rise to the option value of holding the lease as

$$V(q,c) \equiv \max_{\tau \leq T} \mathbb{E}_{P_0, c_0=c} \left[e^{-r\tau} \left((1-\phi) P_\tau q - c_\tau \right)^+ \right]. \tag{2}$$

The value defined above for time $t=0$ depends on the cost evolution over time denoted here as c. To simplify, throughout the model I will assume that drilling costs have the structure $c_t = cP_t^\gamma$; thus, the value of c provides a full description of how costs c_t evolve over time. The value of γ, on the other hand, describes the response of drilling rates to prices: when each barrel of oil is more valuable, providers of the rig collect higher revenues from their services. This effect can be important for the timing of drilling. Intuitively, if $\gamma = 1$, the value of the lease is linear in prices,

[12] A simple extension could account for a lag in dependence of costs on prices.

[13] In practice, wells that begin production tend to extract oil without interruptions — see Anderson, Kellogg, and Salant (2018). With my approach, q in the model can be thought of as some discounted version of the true oil quantity, with discount factor depending on the oil price model, production curves, and time value of money. The key assumption here is that production curves cannot be manipulated.

[14] Kellogg (2014) shows that certain types of investment, namely infill oil well drilling, appear to be optimally timed to the prices of oil.

and optimally timed drilling will happen either immediately or at the very end of the lease. In practice, however, the elasticity of drilling costs in oil prices is not as large, that is, $0 < \gamma < 1$. Yet, the larger the value of γ, the lower the benefit to the lease holders from positive price shocks, and the less incentivized they are to delay the investment.

3.2. Further Assumptions and Properties of the Stopping Model

Throughout the paper, I assume the following about the price process.

Assumption 1. *P_t follows a Geometric Brownian Motion, i.e., $dP_t/P_t = \mu_p\ dt + \sigma_p\ dB_t$. The value of μ_p is such that $r > \mu_p$ where r is the time discount rate of the agents.*

Geometric Brownian Motion, being an exponent of a random walk with drift, is a non-stationary random process. Nonetheless, it has been found to perform well[15] compared to alternative models for prices. The assumption on discount rate is a technicality: when the expected prices show a strong positive drift relative to the discount, the optimal timing of drilling may degenerate into the solution with $\tau = T$. In practice, this assumption is not an issue for oil prices.

Assumption 1 allows for a fairly simple characterization of the drilling decision (see Yam, Yung, and Zhou (2012)). In particular, consider the following problem

$$\max_{0\leq\tau\leq T} \mathbb{E}_{P_0,z_0=z} e^{-rt} \left(P_\tau - z_\tau\right)^+ \tag{3}$$

of picking the stopping time τ given the time discount rate r and investment costs z_τ which depend on prices and thus evolve over time through this dependence. In my model, $z_\tau = c_\tau/[q(1-\phi)]$ simply represents the costs of drilling per unit of output retained by the lease holder. The solution to the problem above can be described by the optimal stopping path $P^*_{t,T}(z)$ for $t = 0, \ldots, T$ such that

$$\tau = \min\{T,\, t : P_t \geq P^*_{t,T}(z)\} \tag{4}$$

where $\tau = T$ is interpreted as no drilling if and only if $P_T < z_T$. In other words, the stopping rule can be formulated as drilling at the first moment when oil prices exceed the threshold described by $P^*_{t,T}(z)$

[15] Postali and Piccheti (2006) find that Geometric Brownian Motion performs on par with more sophisticated models when predicting oil prices 4 to 8 years ahead. Note that all the leases I study involve price prediction for up to 5 years.

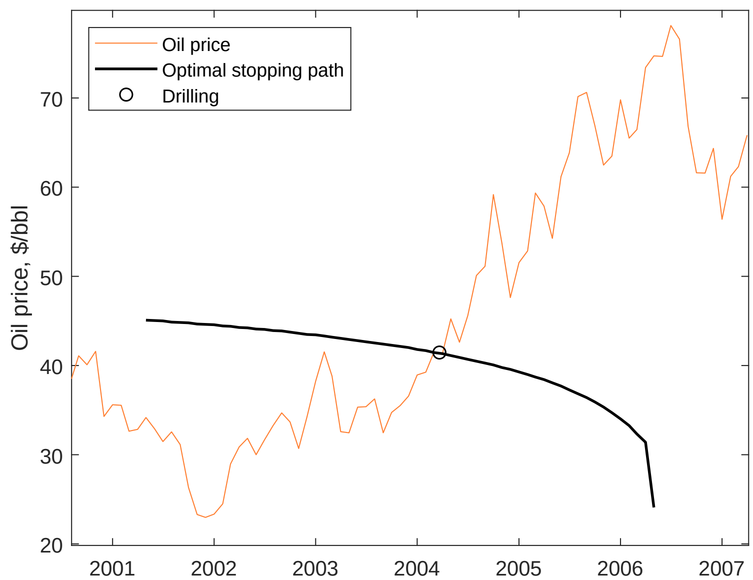

Figure 2: Example of an optimal stopping path $P^*_{t,T}$ based on a well from the data sample.

which depends only on the initial value of costs z. It is convenient to set $P^*_{T,T}(z) = P : P = z_T$, i.e. the threshold price at the last day of the lease is simply the break-even price. Figure 2 illustrates this path for $z_t \approx \$27.5$. Raising z would result in an upward shift for the path, resulting in no drilling for a high enough z.

One implication of this representation is that the optimal stopping path $P^*_{t,T}$ depends only on per-unit costs of the lease holder's output, that is, the $c/(q[1-\phi])$ ratio. For instance, to determine the stopping path in Figure 2, I looked up the lease sale date and drilling date for a well from the data sample and found the value of z such that $P^*_{t,T}(z)$ starts from the sale date at $t = 0$ and intersects the prices of oil precisely at the time of drilling.[16] Only one such path exists due to monotonicity of prices in costs, and it corresponds to the value of z provided above. Further, in this model the distortive nature of revenue taxes involves both the possibility of canceling the drilling decision as well as the deceleration of drilling as the threshold prices $P^*_{t,T}$ increase with ϕ. Another property of this path is monotonicity in t. Intuitively, the more time that remains before the expiration date, the less rushed the lease holder feels to drill and the more interested she is in waiting for higher prices. As time approaches the deadline, however, waiting becomes less attractive,

[16]The stopping paths also depend on the parameters of the price process P_t and the value of γ that measures elasticity of drilling costs in prices. I used the values of these parameters derived in the estimation section below.

and eventually the trigger threshold $P^*_{t,T}$ converges to the breakeven level $P^*_{T,T}$ as the lease time approaches the termination date.

This insight into the optimal drilling problem allows me to further describe the properties of lease value $V(q,c)$ defined in (2). For instance, intuitively enough, $V(q,c)$ is increasing in q and decreasing in ϕ and c. Under Assumption 1, the lease value is strictly increasing in P_0. Furthermore, note that, under the same assumption, the representation

$$\mathbb{E}_{P_0',c}\left[e^{-r\tau}\left((1-\phi)P_\tau q - c_\tau\right)^+\right] = \frac{P_0'}{P_0}\cdot \mathbb{E}_{P_0,c}\left[e^{-r\tau}\left((1-\phi)P_\tau q - c_\tau \cdot \frac{P_0}{P_0'}\right)^+\right]$$

implies that the stopping problem associated with a price P_0' is equivalent to one starting at P_0 but with a different cost level.

Given the characterization of the drilling timing and monotonicity of threshold prices in taxes, it is possible to consider how drilling of the wells observed in the sample would have been delayed under higher ϕ. I address this question in the section that follows.

3.3. Optimal stopping and revenue taxation

Recall that the wells in the data sample for New Mexico are characterized by long delays, with some of them bunching up near lease expiration. As the solution to the optimal stopping implies, these wells must be close to being only marginally profitable; therefore, increases in taxes would phase them out of existence. It does not mean, however, that the state cannot benefit from such policies. If the most productive wells are drilled early and only low-quantity wells are delayed a lot, the state can actually improve its revenues through more aggressive taxation. In this section I let the data establish whether this is, in fact, the case.

First, recall the property of the optimal stopping paths described in Section 3.2 of being monotone in per-unit costs of drilling. Consider now a lease that starts at $t=0$ and gets drilled at time $t=\tau$. Due to the monotonicity property, there is only one optimal stopping path that intersects the observed oil prices at precisely $t=\tau$. Once this path is determined, the well drilled on the lease can be characterized by the value of $c/[q(1-\phi)]$ corresponding to this one path. If the value of ϕ is increased to ϕ', the change in the ratio $c/[q(1-\phi')]$ can be computed and a new stopping path can be determined. This new path would correspond to the same well drilled under a different revenue

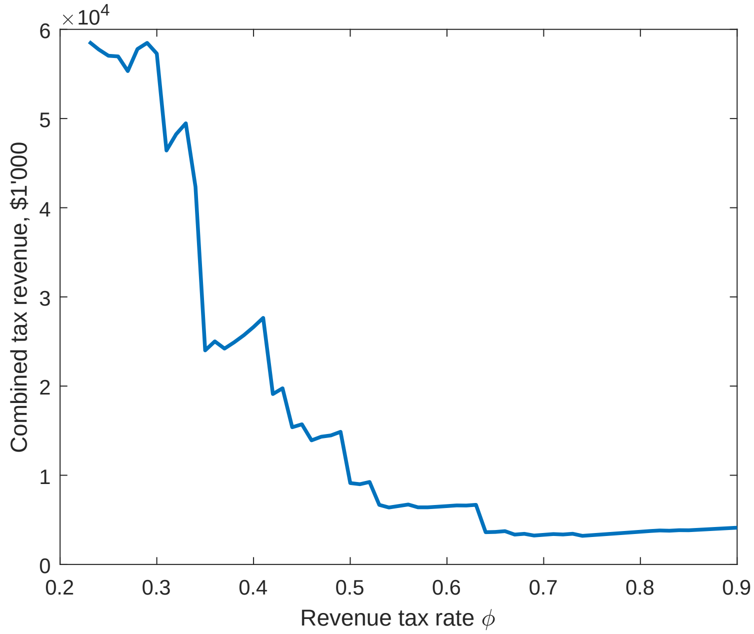

Figure 3: Tax revenue based on the wells in the sample for varying ϕ.

tax. In particular, the new timing of drilling corresponding to ϕ' would be determined based on the intersection of this new path with the observed oil prices. Finally, if q is known from post-auction data, state tax revenues corresponding to ϕ' can be computed based on this method.

I perform this very exercise using the selected data sample with 78 wells. For each drilled lease in the sample, I derive the expected lifetime oil output based on the exponential production decay model (see Appendix C for details). Next, based on the time of drilling, I determine the stopping paths corresponding to these leases. Then I compute

$$R(\phi) = \sum_{i=1}^{Nobs} e^{-d\tau(\phi)} \phi q_i P_{i,\tau(\phi)} \mathbb{1}\left(i \text{ is drilled}\right) \tag{5}$$

where $\mathbb{1}\left(i \text{ is drilled}\right)$ is an indicator for whether lease i would be drilled given the tax ϕ. Note that, starting from the baseline level, I can only increase ϕ here, since this approach does not allow me to account for selection into drilling for leases without wells when ϕ is decreased.

Figure 3 shows how $R(\phi)$ varies. It can be observed that past the threshold of $\phi = 0.29$, which maximizes the revenue, the value of R quickly decreases. Of course, $R(\phi)$ does not include the revenue the state collects in the form of auction bids. Since the bids are decreasing in ϕ, the value of ϕ maximizing $R(\cdot)$ would represent an upper bound on the full revenue maximizing tax rate.

While this exercise is informative of the state's limited potential to raise higher revenues, it does not allow me to account for a variety of effects. Auction participation and bidding are not considered here; furthermore, the extent to which the drilling frequency can increase when distortions are removed cannot be measured with this method. Additionally, more sophisticated policies that could adjust themselves to oil prices or endogenize the tax rates based on the information available to the firms are beyond analysis under this simple approach. Therefore, for further analysis I develop an extension to the model which allows me to investigate these additional policies and their various consequences for the state and for the industry.

4. Model

In this section, I present a model of auctions for real options followed by a dynamic investment problem, and discuss its properties. The model can be viewed as an extension of the contingent payment framework in DeMarzo, Kremer, and Skrzypacz (2005).

4.1. Setup

A set of $N \geq 2$ bidders compete for a tract of land in a first-price sealed-bid auction. There is a quantity q of oil, which is common to all bidders, that can be extracted from the ground. Each bidder gets a signal $q\xi_i$ about the quantity that she can extract, with private noise ξ_i independent of q and $\mathbb{E}[\xi_i] = 1$. Each agent, if deciding to participate, submits a cash bid b_i for the right to drill on the particular tract of land. The bidder who submits the highest bid wins the right to drill on the land for a period of T years.

The auction winner foregoes a payment of X immediately and learns the true value of q as well as the drilling costs c_0 at time $t = 0$. Subsequently, as prices evolve for $t \in [0, T]$, the winner may decide to drill a well. If such a decision is made, production is realized instantaneously. Revenue tax at the rate of ϕ is collected by the auction holder, i.e. the state, and no further drilling is assumed to happen.[17] If no well is ever drilled, the lease expires at $t = T$.

Thus, the model consists of two stages: bidding followed by drilling. The second stage of it is discussed above. One detail that requires clarification here is the the assumption on cost distribution.

[17] In practice, firms sometimes drill multiple wells on the same lease. An extension of the model would allow for drilling of additional wells.

Specifically, I assume that $c_t = cP_t^\gamma$ and the costs c have the same distribution for every agent[18] with a cumulative distribution function given by $H(\cdot, q)$.

4.1.1. The Bidding Stage

The drilling problem from (2) gives origin to bidder values. Let

$$v(q) \equiv \mathbb{E}_{c \sim H(\cdot,q)} V(q, c) \tag{6}$$

denote the ex-ante value for drilling rights at the bidding stage if the true quantity of oil is q. This quantity feeds into the signals of values at the time of bidding since (i) bidders are uncertain about the true quantity as well as their idiosyncratic costs at the time of bidding, and (ii) bidders face the standard "winner's curse" concern that the signal conditional on winning is an overestimate of the true quantity. In particular, participants will choose their bid $b(s_i)$ as a function of their signal $s_i = q\xi_i$ by solving

$$\max_{b \geq 0} \mathbb{E}_q \left[(v(q) - X - b) \cdot \mathbb{1} \left\{ \max_{j \neq i} \beta(s_j) < b \right\} \middle| s_i = q\xi_i \right] \tag{7}$$

where the expectation reflects the dependence of the bid on the private signal as well as the outcome of winning in the auction. In equilibrium, each agent bids as if her signal were marginal (i.e., coincides with the second-highest signal). Agents who expect profits less than X do not bid.

To establish the existence of equilibrium in this model, I impose a number of technical conditions on the quantity and cost distributions.

Assumption 2. *The densities of q and ξ are positive on compact support and $\mathbb{E}[\xi_i] = 1$. The distribution of c as a function of q is such that $\mathbb{E}_{c \sim H(\cdot,q)} V(q, c)$ is increasing in q.*

Assumption 2 ensures that costs rise slowly enough as a function of quantities q so that the expected value from holding a lease still increases in q. If per-unit extraction costs are stochastically decreasing with quantities, then it is easy to check that this assumption is satisfied, as would be the case, for instance, in a model where the fixed and marginal costs of drilling are drawn from distributions that

[18]Cost asymmetries based on observable bidder characteristics would be straightforward to incorporate in the model.

are independent of quantity. This monotonicity assumption also implies that bidders with higher signals of q are also more optimistic about the value of obtaining the rights to drill.

Equilibrium in the bidding stage is delivered by the classic result for auctions in common values detailed in Milgrom and Weber (1982). Indeed, the signal structure at the first stage of the model possesses the affiliation property, and monotonicity of agents' valuations in their signals is guaranteed by the assumption above. Consequently, the bidding strategy in the first stage of the game is described by

$$\beta'(s) = \left[\mathbb{E}_q\left(v(q) - X | s_1 = s, \max\{s_2, \ldots, s_N\} = s\right) - \beta(s)\right] \frac{f(s|s)}{F(s|s)} \tag{8}$$

where $f(x|y)$ is the density of the maximum among $s_2, \ldots, s_N$ evaluated at x conditional on $s_1 = y$, and $F(x|y)$ is the corresponding cumulative distribution function. The initial condition for the bidding is given by $\beta(\underline{s}) = 0$ where $\underline{s}$ solves

$$\mathbb{E}_q\left(v(q) | s_1 = \underline{s}, \max\{s_2, \ldots, s_N\} = \underline{s}\right) = X.$$

Theoretical results for settings with these features typically indicate that a priori it is not clear what level of taxes, or, in fact, what alternative contract design (save for profit sharing) could maximize the revenues of the seller (see, e.g., Cong (2019) or Board (2007) where similar two-stage frameworks are considered). Therefore, understanding the policy effects in this setting calls for an empirical approach. Before estimation, however, I discuss how the implications of the model match up with the data.

4.2. Discussion of Modeling Choices

In the model outlined above, the behavior of the agents significantly depends on the level of prices. In Table 2 I collect empirical evidence towards this and discuss it below.

First, column (1) shows results for a basic regression of bids on the tract size and oil prices while controlling for the fixed effects of sale year and the fixed effects of the number of bidders. Qualitatively, the results are aligned with the model: there is a statistically significant positive relationship between prices and bids. Specifically, every increase in prices by \$1 results in approximately a \$924 rise in

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Bid 1	$\mathbb{1}$(drilled)	$\mathbb{1}$(drilled)	Delay	Log(Bid 1)	$\mathbb{1}$(drilled)	$\mathbb{1}$(drilled)	n/N
P_0	0.924***	0.200**	0.018	7.362***		0.250**	0.338***	0.001
	(0.256)	(0.101)	(0.013)	(1.792)		(0.106)	(0.110)	(0.001)
Acreage	0.154***	-0.003	-0.000	0.840	0.004***	-0.004	-0.013	-0.000
	(0.013)	(0.008)	(0.001)	(0.589)	(0.000)	(0.014)	(0.016)	(0.000)
Bid 1		0.101***	0.009***	-3.009***		0.017		
		(0.022)	(0.002)	(0.812)		(0.031)		
Q(oil)				0.001				
				(0.001)				
Log(P_0)					0.525***			
					(0.170)			
Bid 2						0.195***		
						(0.053)		
Log(Bid 1)							1.716	
							(2.062)	
Log(Bid 2)							7.250***	
							(2.127)	
intercept	-43.543***	-6.016*	-6.569***	1011.180***	-0.192	7.753	-11.657	0.960***
	(8.726)	(3.245)	(1.442)	(219.886)	(0.558)	(8.610)	(8.381)	(0.040)
N	786	786	757	54	786	506	506	786
R^2	0.420	0.105		0.218	0.538	0.151	0.130	0.228

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2: Basic regressions for sample data. All regressions include year fixed effects, and regressions for the winning bid also include fixed effects for the number of bidders. (1) is OLS for the winning bid, (2) is linear probability regression for $\mathbb{1}$(is drilled), (3) is a logit regression for the same variable, (4) is OLS for the delays. Column (5) recasts the first column as OLS for log-winning bid, (6) is a linear probability for $\mathbb{1}$(drilled), and (7) is a version of (6) with log-bids rather than bids on the right-hand side. Finally, (8) is an OLS for the participation ratio n/N. Throughout, bids are measured in \$1'000, delay is measured in days, and $\mathbb{1}$(drilled) takes values either 0 or 100, i.e. it corresponds to % probability of drilling (except for logit where its values are either 0 or 1). For regressions with the runner-up bid, only auctions with at least two bidders were used (qualitatively, the results are the same if Bid 2 is set instead to 0 when absent). The number of observations in (3) is low because observations for years without drilling were dropped.

the winner's bid. Column (5) establishes a similar result for log-bid of the winner and log-prices, finding that about 53% of changes in the price at sale passes through to the bid.

Other regressions, namely columns (2-4) and (6-7), are dedicated to measuring the responsiveness of drilling decisions to prices and bids. Column (2) — a linear probability model for the decision to drill a well — reveals a positive dependence of the drilling probability on both prices and bids. Notably, this dependence is somewhat weak: the coefficient on P_0 is statistically significant only at the 5% level, and its value implies that a \$10 price increase at auction would lead to a two percentage point boost for the drilling chance. Two comments are due here. First, given the sample average of a 0.1 drilling frequency, these percentage points translate into a 20% increase in the frequency of drilling relative to the sample mean. Second, the model described above demonstrates that P_0 may not necessarily serve a good proxy for incentives to drill. Specifically, the lease holder solves the optimal stopping problem after the auction and over 5 years. During such a long time segment, prices are bound to change a lot. The limitation due to the sample size suggests the possibility that a simple regression cannot properly integrate out all such post-auction variation. A structural model, on the other hand, could properly handle the exact price paths observed in the data. Thus, although the relationship between at-auction prices P_0 and the probability of drilling may be found to lack statistical significance in the regressions, I do not view this as a contradiction with the model[19]. And, indeed, the coefficient on P_0 is insignificant in column (3), which contains the Logit version of the regression in (2). Another finding from these two columns is that the drilling decision is positively related to the winner's bid. In particular, the OLS results suggest that a \$10,000 higher bid, other things being equal, is associated with a 1 percentage point higher chance of drilling. Although the Logit results do not allow for such a direct interpretation, they are qualitatively similar. Computing the marginal effects at the mean, the result is similar: a \$10,000 higher bid is associated with a 0.7 percentage point higher chance of drilling.

The analysis of drilling delays is more complex, since delays are only defined in cases when a well is actually drilled. Additionally, the discussion above about dependence on price paths as opposed to P_0 applies to delays as well. Column (4) makes an observation seemingly contradictory to the model: higher initial prices lead to greater delays. Since other regressions indicate that a higher fraction of leases end up drilled at higher prices, the overall effect for delays could be positive with

[19]Furthermore, regressions where yearly fixed effects are dropped show stronger dependence between the variables.

no contradiction. Dependence on the winner's bid, on the other hand, is in agreement with the theory: a \$10,000 higher winning bid is associated with a delay one month shorter. Admittedly, relative to the typical delay of almost 4 years, this one month is not very impactful.

The setting with post-auction strategic actions allows me to test for the common values assumption. Traditional tests, such as the one presented in Haile, Hong, and Shum (2003), rely on bid data alone. On the other extreme, Hendricks, Pinkse, and Porter (2003) leverage the data that describes both bidding and true valuations of the auctioned items (which are, incidentally, leases for offshore oil territories) to study whether the strategies of the agents are in accord with the traditional models of bidding in auctions. While the post-auction data I use does not allow me to directly evaluate the option value of holding a lease, the decision to drill a well in the common values setting would be dependent on bids of participants other than the winner. Simple regressions allow me to confirm that this is, indeed, the case in the data.

Columns (6) and (7) show that post-auction decisions of the winners do not contradict the common values paradigm of the model. Concretely, both columns provide OLS model estimates for the probability of drilling given the value of the second-largest bid in the auction. In private values, this probability would depend only on the winning bid; regressions, however, highlight that the drilling decision is strongly and positively correlated with the runner-up bid. This conclusion can still be made if regressions are run to account for additional geographic fixed effects or if, alternatively, yearly fixed effects are removed. Similar results hold for the delay variable in regressions analogous to (5); however, given the selection problem, I omit them here. Finally, the last column looks at the participation rate in the auctions. Consistent with the model, the finding is that agents bid more frequently when oil prices are higher: a 5 percentage point rise in participation after one standard deviation increase in prices. Relative to the sample mean, this is about 8% higher participation rate.

An interesting feature of the optimal stopping setting studied in this paper is that the optimal stopping paths discussed in Section 3.2 depend only on the model for prices. Thus, given such a model, it can be tested for each drilling observation in the data whether the model predictions about the timing are consistent with it. In Appendix D I perform a basic analysis which shows that the majority of drilling events from the sample could be described as solutions to the optimal stopping problem based on the model considered above.

5. Estimation and Results

The revenue taxes described in Section 2 are used in estimation and represent ϕ. On the other hand, the quantities and the costs that are recovered in estimation represent the values *perceived* by the firms. Thus, they may absorb various effects such as time discounting or the deduction policies similar to the intangible costs rule that the firms recognize.

Identification results for a version of the model where costs $c_{i,t}$ do not vary with prices have been established in Bhattacharya, Ordin, and Roberts (2019). For the purposes of estimation below I evaluate the dependence of costs on prices in a reduced-form manner[20] using data external to the model itself.

5.1. Estimation

My estimation approach is applied under a common values paradigm and leverages parametric assumptions. I use a method of simulated moments approach to estimation that proceeds in four steps. First, I estimate the parameters of the price process. This price process can give me the optimal drilling path for any quantity q and initial cost c. Second, I derive the parameter governing the dependence of costs on P_t using the data on the costs for primary services in oil and gas. Third, for each set of candidate parameters for cost and quantity distributions, I can simulate the bidding and drilling decisions of the agents. Fourth, I then match a set of model-implied moments to the data.

5.1.1. Price Process Parameters

Under Assumption 1 on the price process,

$$P_t = P_0 \cdot \exp\left[\left(\mu_p - \frac{\sigma_p^2}{2}\right)t + \sigma_p B_t\right]. \tag{9}$$

The price process is observed in discrete intervals $\Delta = 1/12$ corresponding to one month. Taking logs of (9) gives

$$\log\left(P_{t+\Delta}\right) - \log(P_t) = \left(\mu_p - \frac{\sigma_p^2}{2}\right)\Delta + \sigma_p\sqrt{\Delta}\epsilon_t, \tag{10}$$

[20] I have also estimated the parameters of interest without this reduced-form approach, resulting in an estimate close to the one found below.

	Lag0	Lag1	Lag2	Lag3	Lag4	Lag5	Lag6	Lag7	Lag8
log(P)	0.27**	0.28**	0.30**	0.31**	0.31**	0.31**	0.31**	0.32**	0.32**
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
N	61	61	61	61	61	61	61	61	61
R^2	0.84	0.87	0.91	0.93	0.94	0.93	0.93	0.92	0.90

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$

Table 3: Regression for log-cost of well drilling services in oil and gas industry on lagged log-prices of oil. Lags are in months. The data for the costs spans Jan 2014 through Jan 2019.

where ϵ_t is the Brownian innovation, and thus ϵ_t are i.i.d. standard normal. Parameters of interest in (10) are estimated via Maximum Likelihood (see Appendix B.2).

5.1.2. Price Component of Costs

To model the dependence of costs on contemporaneous prices, I assume the following structure for each bidder i:

$$c_{i,t} = c_i P_t^{\gamma} \tag{11}$$

The parameter of interest in this equation is γ; essentially, it measures the passthrough of oil prices into the well-drilling costs. For simplicity, the dependence across time as modeled here is without lag, even though in practice the upstream market for a rig may respond to price shocks with a delay of several months. To show how this assumption squares up with the data, I run a series of regressions using a proxy for log-costs as a dependent variable and log-prices with different lags as regressors. My proxy for drilling costs is the index for cost of drilling services in the oil and gas industry in the US.

Table 3 reports the estimation results. For a broad range of lags in prices, the coefficient γ remains approximately the same: around 0.3. A similar finding is made in Osmundsen, Rosendahl, and Skjerpen (2015), where the same passthrough level of prices into costs is estimated using the data on offshore well drilling costs. In fact, even after developing a model with capacity constraints for the upstream market, the authors find a similar passthrough level. Consequently, in my estimation I simply set $\gamma = 0.3$.[21]

[21] When estimating γ together with other parameters using the method of simulated moments, I find the minimized function to have low sensitivity in γ and arrive at the estimate $\hat{\gamma} = 0.32$ after using a set of different values as the

5.1.3. Optimal Stopping Trajectories

The next step involves computing the drilling rule $P^*_{t,T}(z)$ as a function of c and (implicitly) q. Specifically, the goal is to solve

$$\max_{\tau} \mathbb{E}_{P_0,z_0=z} e^{-rt} \left(P_\tau - z_\tau\right)^+ \tag{12}$$

where z_τ starts at level $z_0 = z$ and evolves over time, broadly speaking, according to some law that depends on the prices P_t contemporaneously. In the problem above, z corresponds to the ratio $c/[(1-\phi)q]$ of the general model. If the optimal stopping rules can be found for the problem of the type (12), they can be used to describe the solution to the stopping problem for lease holders.

I characterize the solution to (12) with the path $P^*_{t,T}(z)$, which is defined as the minimal price which at time t induces the decision to drill given that the cost of investment at $t = 0$ was equal to z. For simplicity, I discretize time t at the monthly level as if the lease holder could make the decision to drill only once each month rather than at any point in time. Quantitatively, such an approximation ends up close to the continuous time solution. Note that it is the properties of Geometric Brownian Motion and the assumptions about the way in which z depends on P_t that allow such monotone characterization in the first place. Additionally, it will be useful to introduce valuations $w(t, P, z)$ corresponding to the value of holding lease rights *at time* t given price P and (initial) investment costs z.

Under the assumption that the agents make decisions at a monthly frequency and given the finite time horizon of T years, the problem can be solved from the end. First, at $t = T$ note that drilling happens if and only if $P_T \geq z_T$; thus, $P^*_{T,T}(z) = z_T$. Consequently, $w(T, P, z) = \max\{0, P - z_T\}$. Going back in time, in month t the continuation value is described by

$$\begin{gathered} w(t, P, z) = \max\left\{P - z_t, e^{-r\Delta}\mathbb{E}w(t+1, P', z)\right\} \\ P' = P\exp\left[(\mu_p - \sigma_p^2/2)\Delta + \sigma_p\varepsilon\right] \end{gathered} \tag{13}$$

where $\Delta = 1/12$ is one month and the expectation is taken over the Brownian innovation ε which determines the evolution of $P \to P'$ and, consequently, the implicit evolution of z_t. The maximum

starting points for optimization.

on the right-hand side represents the choice bewteen two options: drill now or wait until later. The optimal stopping rule $P^*_{t,T}(z)$ is then delivered by

$$P^*_{t,T}(z) = \min\{P : \ w(t, P, z) = P - z_t\} \tag{14}$$

Computationally, given the price process parameter estimates μ_p and σ_p, for a range of various values of z equation (13) is solved from the starting condition $w(T, P, z) = \max\{0, P - z_T\}$. The burden of these derivations would be greater if the valuations w depended not just on P but also on, say, the volatility of P_t. The Geometric Brownian Motion assumption allows me to solve the equation much quicker. Finally, $w(t, P, z)$ computed, the optimal stopping rules are derived directly from (14).

While the optimal stopping paths derived here are used to determine the timing of drilling, the normalized valuations $w(0, P, z)$ are used to derive the bidder valuations. Specifically,

$$V(q, c) = w\left(0, 1, \frac{c}{(1-\phi)qP_0}\right)(1-\phi)qP_0 \tag{15}$$

where $V(q, c)$ is the valuation from (2).

5.1.4. Parameterization

The rest of the estimation relies on more exact parametric assumptions. For a given auction,

$$\begin{aligned} q &\sim Exp(\lambda_q), \\ \xi_i &\sim Weibull(\lambda_{\xi,0}, \lambda_{\xi,1}), \text{and} \\ c_{i,t} &\sim \log\text{-}\mathcal{N}\left(\mu_0 + \alpha \log(q) + \gamma \log(P_t), \sigma_c\right). \end{aligned} \tag{16}$$

Thus, the common component of quantities q, which can be intepreted as a measure of the *true* quantity of oil, is distributed exponentially with mean parameter λ_q. The idiosyncratic shocks ξ_i have a Weibull distribution with scale and shape parameters $\lambda_{\xi,0}$ and $\lambda_{\xi,1}$. The parameters are normalized so that $\mathbb{E}\xi_i = 1$. Finally, I have $\log c_t$ distributed normally with mean parameter $\mu_c = \mu_0 + \alpha \log(q) + \gamma \log(P_t)$ and scale parameter σ_c. Over time, costs change according to $c_{i,t} = c_{i,0}(P_t/P_0)^\gamma$.

5.1.5. Method of Simulated Moments

Given a candidate set of parameters, auction-specific N and P_0 observed in the data, and some draw of q and agent-specific shocks ξ_i, valuations $v(q)$ from (6) are derived based on the computations described in Section 5.1.3. Agent signals $q\xi_i$ are then mapped into bids as described in Section 4.1.1 and equation (8). Finally, with some draw of winner's costs $c_{i,0}$ and a price path originating at P_0, I compute the solution to the optimal stopping problem, deriving the delay from sale month for cases when a well is drilled. For every auction in the sample, multiple draws, namely 100, are taken.

Based on these simulations, I compute bidding- and drilling-related moments. On the auction side, I compute the winning bid, the runner-up bid, the average bid across all agents, and the number of bidders. On the post-auction side, I compute the probability of no drilling, the probability of immediate drilling, the average at-auction prices for auctions not followed by drilling, and delays. I make sure to avoid working with conditional moments. For instance, when in simulations there is only one participant in the auction, I interpret it as runner-up bid being equal to 0. In cases of no drilling, I set the delay value to the largest possible number: 1,825 days, corresponding to the full 5 years. Although my approach allows me to also account for moments based on the quantity of oil extracted, I avoid doing this given that in the data observations for q are sometimes missing.

It is not unusual to rely on simulation methods to estimate a bidding model (see, e.g., Krasnokutskaya and Seim (2011)). The reason for my not using a likelihood method has to do with outliers in bids. In an auction model, it is a typical situation to have the bids cap off at some finite level even when signals and valuations of bidders have infinite support. If for a given draw of parameters at least one observed bid exceeds this level, the likelihood is not defined at all. Thus, to freely navigate the parameter space during estimation, I follow the simulated methods approach where this feature of the model is not a problem.

5.2. Results

The parameter estimates are reported in Table 4 and the model fit in Table 5. As some of the parameters are difficult to interpret on their own, I instead point out what they imply. The standard deviation of the noise ξ is estimated to be 1.34, which corresponds to a within-auction correlation of signals of about 0.2 and a correlation of 0.46 between q and a draw of $q\xi_i$. Thus, the bidders'

λ_q (K bbl)	StDev of ξ	μ_0	α	σ_c	X($)
76.5	1.42	6.420	0.756	0.875	6,886
(0.05)	(0.004)	(0.04)	(0.005)	(0.01)	(40)

Table 4: Parameter estimates. Standard errors are based on estimates for 50 boostrapped samples.

information about the true quantity of oil at the time of the auction is rather imprecise.

The average hypothetical well is expected to produce about 76,500 barrels of oil. However, in simulations, wells which are drilled are expected to produce 114,600 barrels. Compared to the average of 75,400 barrels produced by all the wells in the sample for which production data is available, this is fairly high. The model, however, does not take into account the possibility of multiple wells being built on the same tract. Of course, the estimation procedure does not use the production data, and this comparison represents something like an out-of-sample test of the model.[22]

Cost parameter estimates are not straightforward to interpret either. For instance, for the typical values of q (set to the expected output conditional on drilling) and P_0 (set to \$50), the average costs end up being about \$15mln — a rather large value. The reason for this is simple: in the simulations, just like in the data, drilling is very infrequent. The high value of costs is the way the model explains the lack of drilling. To understand what the cost estimates imply, I compute the drilling costs in simulations for the cases when a well is drilled. The average costs across the sample end up being about \$3.1mln, a far more reasonable value given the time range of 1994 through 2012. Furthermore, costs conditional on drilling show moderate variation. For instance, their 95th percentile is \$5.9mln.[23] In addition, it should be stressed that these are real costs as perceived by the lease holders. The intangible deduction alone can reduce the nominal costs by up to 30%; thus, it is not surprising that these cost levels may be somewhat lower than the data in publicly available reports.

[22]One could also view this discrepancy as a consequence of q absorbing the possibility of multiple well drilling. Denoting the value of drilling extra wells when the oil prices are at the level P as $V(P)$, the driller's problem could be represented as (cf. (2))

$$\max_{\tau} \mathbb{E}_{P_0,c_0=c}\left[e^{-r\tau}\left((1-\phi)P_\tau q + V(P_\tau) - c_\tau\right)^+\right]$$

with the $V(P_\tau)$ term reflecting the additional incentive to drill. Thus, the value of q conditional on drilling could be exaggerated in the estimation results due to the bias arising from this extra term. In further work, I aim to account for this possibility of multiple wells.

[23]This is close to the typical drilling costs for various regions in 2014, see, e.g., Energy Information Administration (2016).

	Winning Bid ($1'000)	Avg Bid ($1'000)	P(Drilling)	Delay (days)	Q oil (1'000 bbl)
Model	61.4	25.8	0.096	1'253	114.6
Data	65.9	27.0	0.099	1'275	75.4

Table 5: Model fit. Moments in simulations and in the data at the estimated parameter values.

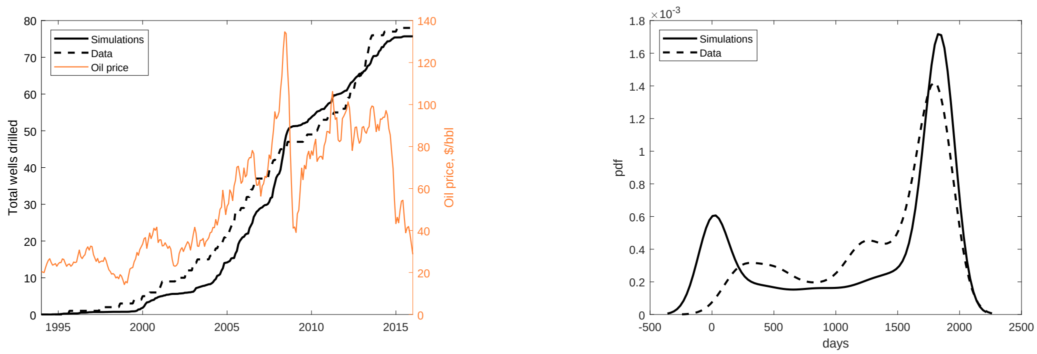

Figure 4: Simulations and data: wells drilled over time and distribution of drilling delays.

Finally, the winner's payment of X is estimated to be $6,886. This estimate includes any costs associated with owning the lease that are independent of drilling, such as managing the paperwork or land rents paid to the NMSLO. For reference, the typical annual land rents of $1 per acre would contribute $1,600 to this payment for the standard 320 acre plot of land leased for 5 years.

As Table 5 shows, the model is doing fairly well at matching the moments of interest to the sample averages. In particular, the winning and the average bids, though slightly understated, are very close in simulations and in data. The fit is very close for the post-auction outcomes, both for the probability of drilling and the average delays.

Finally, as an additional measure of fit, I describe the drilling decisions of the auction winners in more detail in Figure 4. The figure on the left shows when wells are being drilled over time in the data and in the simulations. The plot indicates that the model predictions closely follow the data across the time, even though the estimation procedure involves minimizing the distance between the two only at one point: the termination date of the last lease in the sample. The right figure, showing the distribution of delays between auctions and well drilling, demonstrates that the tendency of firms to drill at the very end of the lease is also a feature of the optimal stopping simulations.

6. Counterfactual Policy Evaluation

The current practice of imposing revenue taxes on oil and gas wells around the US may have significant distortionary effects. As illustrated by the model, well drilling can be delayed if not cancelled altogether as a consequence of these taxes. Given the model estimates, the purpose of this section is to evaluate the distortions and determine whether the state can benefit from policies other than revenue taxation.

I consider three groups of interventions: cost-side, revenue-side, and profit taxes. Another idea I explore relates to the problem of mechanism design in the leasing setting. First-price auctions commonly used for the public lands feature bidding that is linked only to ex-ante information available to the bidders. Can the ex-post realizations of costs, quantities, and prices of oil be used instead? Profit taxation is a tool that asserts observability of all these variables, but finds limited application for various reasons.[24] I look at two alternatives which leverage the data that is already observed anyway: sliding scale royalty and equity bidding.

For the simulations below, I am working with a single auction with fixed features. These features are chosen so that the results are qualitatively and quantitatively similar to those obtained by performing analogous simulations for every single auction in the data sample. To predict post-auction outcomes, I simulate 10,000 price paths and take the average of the variables of interest across these paths. These simulations are performed according to the model for the price process estimated in Section 5.1.1. The auction itself is simulated 20,000 times to integrate out the joint distribution of q and c, as well as the signals received by the agents before the bidding stage.

Quantitatively, I find that only somewhat minor improvements in state revenues are possible through simple policies, which is consistent with the results established in Section 3.3. The more sophisticated taxation systems, however, can have strong effects. A sliding scale royalty can raise the revenues of the state without significantly distorting the drilling decisions; on the other hand, equity bidding depresses economic activity dramatically, reducing the state's revenues by about 50%.

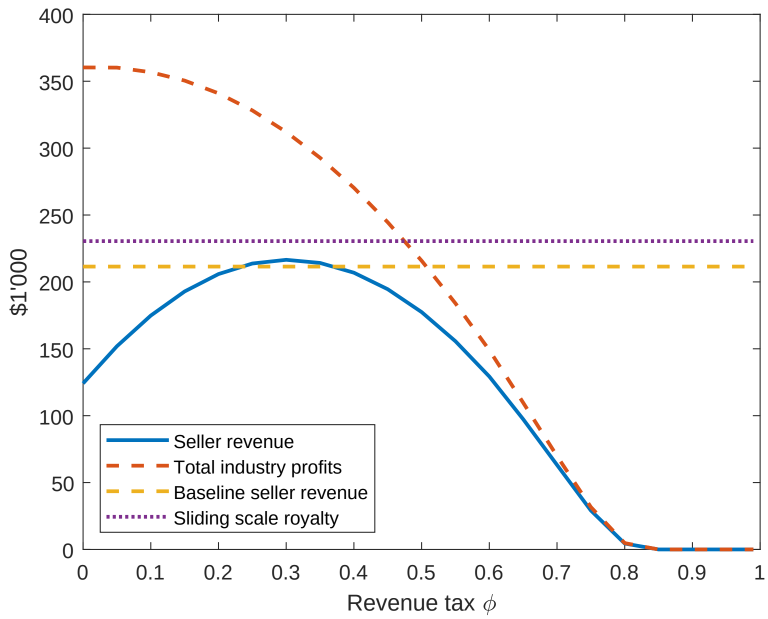

Figure 5: Expected payoffs for the seller. Total profits include the seller's revenue and the winner's profits.

6.1. Revenue taxation

Perhaps the most typical approach in the US to collecting revenues from the oil and gas industry is production taxes and royalties.[25] Occasionally, states review and redesign their severance taxes (though, typically, the magnitudes of these changes are rather low). The purpose of this section is to understand the potential impact of such policies for New Mexico. I consider alternative royalty rates as well as the sliding scale royalty system.

Figure 5 presents the results. Simple taxation alternatives seem only moderately improve the state revenues. The drilling activity, however, can be majorly improved if the revenue taxes are lifted: the probability of drilling goes up from 0.096 to 0.154 as the total industry profits only increase by about 8%. Considering that the state likely cares about the overall level of economic activity, the current tax could in fact be close to optimal from the regulator's perspective.

A sliding scale royalty is a revenue tax with a rate that is set according to the level of the underlying commodity price. For instance, Saudi Arabia's Aramco faces a royalty rate anywhere

[24] Researchers bring up problems such as imperfect monitoring and firm risk aversion when explaining why profit contracts are not used more frequently — see Robinson (1984) for a discussion.

[25] According to the New Mexico State Land Office reports, for the fiscal year of 2017 between bids and royalties about 15% of the office's revenue came from the cash bids. As another example, for federal lands the distribution is skewed even further towards the tax revenues.

between 20% and 50% depending on the prices of oil. The windfall profit tax introduced in the US in the early 1980s (Rao (2018)) was implemented in response to the sudden oil price spike that occured in 1979. When prices fell back down, the tax was removed.[26] Similar to these examples, I consider a revenue tax ϕ that changes depending on P_t. The more involved sliding scale royalty can result in a meaningful boost to revenues. I search through the schemes described by

$$\phi(P_t) = \phi_0 + \mathbb{1}\left\{P_t \geq \overline{P}\right\}(P_t - \overline{P})\phi_p \tag{17}$$

for various values of $\overline{P}$ ranging from \$25 to \$100, ϕ_0 between 0.1 and 0.3, and ϕ_p between 0.0005 and 0.005. I find that the revenue maximizing combination corresponds to $\overline{P} = \$50$, $\phi_0 = 0.225$, and $\phi_p = 0.001$.[27] This implies rather mild variation in taxes; for instance, for prices rising from \$50 to \$100 the tax would grow from 22.5% to 27.5%. As a consequence, however, state revenues are increased by almost 10% relative to the baseline.

6.2. Equity bidding

While the standard royalty rates ensure that the state's payoff is proportional to the sale prices and extracted quantities, the rates themselves are not readjusted.[28] Here, I consider an auction that allows the firms to bid in tax rates instead of cash, i.e. an equity auction. The second stage remains the same: the highest bidder gets to hold the lease over the finite time period. This idea is in part inspired by the use of equity bidding in other settings such as company acquisition as well as the practice of negotiating over royalties in private leasing on US lands.

The model describing this competition is fairly similar to the baseline. One major alteration is the bidder's maximization problem. Instead of (7), the strategy is the solution to

$$\max_{b \geq 0} \mathbb{E}_q \left[\left(v(q[1-b]) - X\right) \cdot \mathbb{1}\left\{\max_{j \neq i} \beta(s_j) < b\right\} \middle| s_i = q\xi_i \right]. \tag{18}$$

[26] In some states, mild versions of the sliding scale are used today. For instance, severance tax rates are set in North Dakota depending on the previous year's average commodity prices. The range of variation, however, is rather low: between 5% and 6%. International policies in the upstream oil and gas provide more dramatic examples.

[27] This form for the sliding scale is inspired by the approaches followed in the real world, where, starting from some initial level, royalty gets higher every time price reaches one of the pre-established thresholds. For simplicity, I look at steadily rising royalties rather than threshold rules.

[28] That is, not readjusted in the US. Over the world royalties are sometimes set based on productivity of wells. For instance, up north in Alberta (Canada) the typical royalties range from 15% to 40% depending on average daily output of the well.

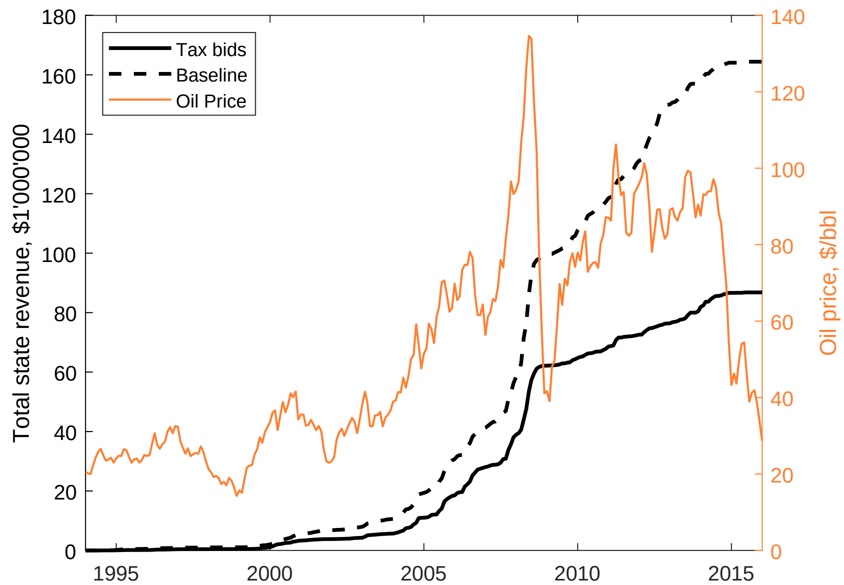

Figure 6: Expected payoffs for the state accumulated over time from equity bidding and baseline bonus bidding. Simulations are based on the real prices of oil and all leases in the data.

Such bids are naturally adapting themselves to the information available to the bidders before the auction. In the common values setting, however, the winner's curse problem is intensified when using equity bidding: the most optimistic firm may submit a bid so high as to discourage own investment following the auction.

Counterfactuals for a single auction with $P_0 = \$50$ and $N = 3$ show that the average winning bid is equal to about 0.4, which is far more distortive for the drilling decision relative to the baseline of $\phi = 0.23$. As a result, both the probability of drilling and the state's payoff are roughly halved. Since this outcome may be a consequence of the assumptions about the auction, I perform equivalent simulations using the observed oil prices and auctions in the sample. Although the average bid across all the auctions is about 0.28, the state's payoff is still dramatically low as opposed to the result in Section 6.1 just above. This is because the mean winning bid shows a lot of variation over time depending on the price level, the intensity of competition, and the winner's signal at the auction. As Figure 6 illustrates, seller's revenues are consistently lower under equity rather than the standard bonus bidding. As oil prices go up, so do the bids, persistently undermining the drilling incentives of the firms.

Theoretical results and simulations (e.g. Cong (2019)) suggest that it is an empirical question whether equity outperforms the standard fixed-royalty bidding. What is it about the setting considered here that leads to the conclusion above? The intuition is fairly straightforward: firms do not incur particularly great losses if they do not drill at all. Consequently, competition drives them to put up such high bids that drilling only happens when the oil prices following the auction go up within the time span of the lease.

6.3. Cost-side intervention

The purpose of this exercise is to evaluate the consequences of shifting the drilling costs of the agents. I consider the situation where the costs change as

$$cP_t^{\gamma} = c_t \to \alpha c_t \tag{19}$$

so that the seller subsidizes $(1 - \alpha)c_t$ worth of drilling costs in cases when the well gets constructed. When $\alpha = 1$, there are no changes; this represents the baseline case. When $\alpha > 1$, the seller effectively imposes additional costs on the lease holders and subsidizes them when $\alpha < 1$. In the real world, different pieces of regulation act in opposite directions. Cleanup rules, for instance, indicate that the firm must prevent potential well leaks, which effectively means higher costs of drilling. On the other hand, tax rules such as the intangible cost deduction offer generous cost deduction policies to the firms, thus making drilling cheaper for them. For the simulations studied here, α measures the deviations from the status quo in the state of New Mexico.

In addition to changing the drilling decisions, the policy described here affects the bidding strategies of the agents. In other words, both reductions and increases in perceived costs pass through to the auction stage. My model allows me to take this effect into consideration, which is important for measuring both revenues and drilling.

To proceed, I simulate an auction with a specific set of attributes. In particular, I set the number of potential bidders to the sample median: $N = 3$. Initial price P_0 is set to \$50 and a collection of random price paths is simulated starting from this position according to Assumption 1. The taxes are set according to sample averages, so that the revenue tax faced by the producers is 23%, while the seller also collects an additional 9% profit tax. For this auction, I compute the average winning bid, the average profits of the winner and the seller, and the expected drilling rates.

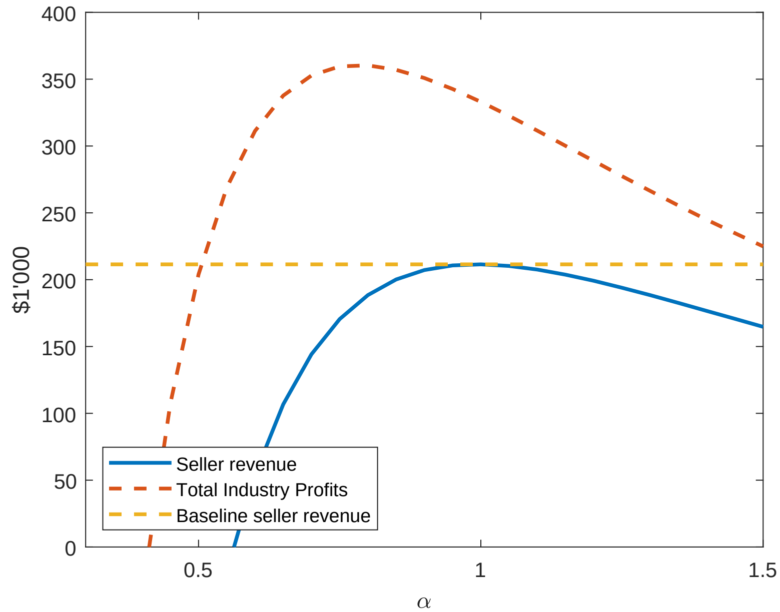

Figure 7: Expected payoffs for the seller. The seller subsidizes $(1-\alpha)$ fraction of drilling costs.

The main result is illustrated in Figure 7. The first conclusion it provides is that cost-side policy cannot substantially improve the revenues of the state. On the other hand, drilling activity can be successfully stimulated through subsidies. Looking at the total industry profits, relative to the baseline of $\alpha = 1$, the tax-cancelling subsidy of $\alpha = 1 - 0.23 = 0.77$ improves the total profits by roughly 8%. This amount of 8% represents the distortionary effect of the taxes currently implemented in the state. While the state is likely interested in both the economic activity and the revenues it collects, the figure suggests that revenue incentives alone would not push the state to implement such subsidies.

To understand the channels through which the policy changes the revenues, it is useful to take a closer look at the drilling rates and the winning bids. Figure 8 shows the effects on the two variables of the cost-side policy considered here. On the left plot, the winning bid is compared to the subsidy that the seller has to give up. Positive values of the subsidy indicate the seller's loss, while negative values correspond to taxation. For the baseline case, $\alpha = 1$, the subsidy is exactly 0. The plot highlights how weak the pass-through effect is: only up to roughly 1/4 of the subsidy is returned to the seller through the auction for a low α. The rest comes from the taxes on production as more wells get drilled. The right plot shows the rise in the probability of drilling. Interestingly, though

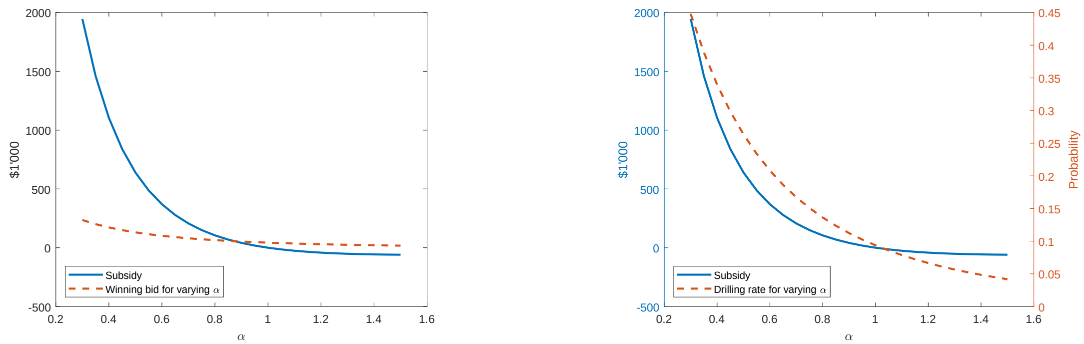

Figure 8: Bidding pass-through and drilling rates under varying cost-side policy.

the optimal (in terms of total industry profits) subsidy would improve the total profits only by 8%, it would increase the probability of drilling by almost 60%. In other words, the model suggests that there are many undrilled marginally unprofitable wells in the baseline state. Such wells are characterized by high costs per unit of output. Elasticity of the drilling probability is even higher for a lower α as suggested by the figure. That is, further cost subsidies would elicit an even stronger response on the intensive margin.

6.4. Profit sharing

As the Norway Ministry of Petroleum and Energy indicates, the petroleum taxation system of the country is intended to be neutral: investment projects profitable pre-tax should remain profitable post-tax. In 2006 Norway collected royalty taxes for the last time, relying instead on the industry-specific 55% profit tax. More broadly, for production service agreements between various governments and contractor firms it is common to include profit-sharing rather than revenue-sharing clauses. I evaluate the potential benefits from using such contracts below.

For the purposes of this section, I still consider firms that compete in first-price sealed bid cash auctions for the leases. These leases, however, have royalties and production taxes replaced with profit taxes instead. Specifically, after paying the federal-level corporate tax, the driller hands over a fraction of the remaining profits to the state. These terms can still introduce a distortion that reduces the participation rates of the firms (I assume here that the payment X is not contractible), but, conditional on participation, the investment decisions of the agents are efficient.

Figure 9 reports the results. It shows that the state can set up a very high profit tax, maximizing

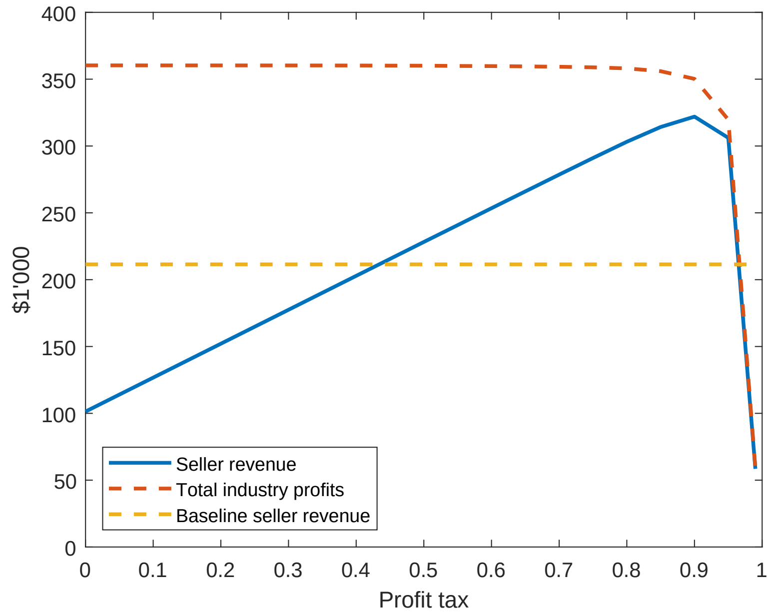

Figure 9: Expected payoffs for the seller and industry profits under profit sharing.

its own revenues at the 90% level. Beyond that point the participation effect becomes strong enough so that firms show interest in the contracts only when they expect particularly high quantities of oil to be extracted. In other words, these contracts can bring the government very close to the most efficient case where the landowner executes the lease option on his own.

This simulation is useful in several ways. While it is not surprising that profit taxation performs well, it is valuable to see that a) the participation margin starts to make a significant impact on the industry activity only at very high profit sharing levels and b) the revenue-equivalent state profit share which would alleviate the distortions from revenue taxes is, as suggested by Figure 9, just above 40%. This number is, in fact, comparable to the 35% profit tax currently established in the state of Alaska.

7. Conclusion

Production contracts with clauses that inefficiently distort the investment decisions of firms are commonly used around the world. On the international level, the oil industry I study here delivers examples of extreme variation in arrangements for extraction of the mineral resources.[29] At the

[29] Consider, for instance, the case of Bolivia where in 2006 president Evo Morales ordered the military to occupy the gas fields of the country and put forward an ultimatum for the firms operating the fields: redirect 50% of the

same time, the industry can be so important that some countries end up fully overhauling their tax treatment of upstream oil and gas in response to international commodity price shocks. The US takes a markedly different approach: though for many states oil and gas each year contribute billions of tax dollars to the government coffers, and though proposals are made to readjust the tax systems in place (Bureau of Land Management (2015)), few policies from these proposals are realized. This paper aims to develop the tools for empirical analysis of production distortions in environments without tax variation and evaluate the potential effects of alternative taxation policies.

A convenient feature of land leases for mineral resources is that the profitability of potential investment can be evaluated even if exact data on drilling costs and oil production is unavailable. My approach relies instead on mapping strategic decisions of the firms, including their bids in land auctions and their choice of investment timing, into the primitives of interest. This way I can predict how various potentially highly complex policies can change the profitability of the wells and thus the incentives to drill them.

My results indicate a limited scope for improving state revenues: the strongest increase of 10% is achieved with a tax rate that is readjusted depending on the prices of oil. On the other hand, the impact of taxes on producers is quite significant: removing the currently existing distortions would raise the drilling frequency by 60%. Looking closer at the features of the data that lead me to these conclusions, I note that many firms in the sample appear to drill wells that seem only slightly profitable. Whether such wells constitute a worthwhile investment can therefore strongly depend on the imposed taxes.

Compared to revenue sharing, profit taxes in this setting would do much better. In this paper, such taxes can be viewed as a reference for the state's ability to absorb industry profits without undermining the investment decisions of the firms. Why are these profit sharing agreements, well-known for their non-distortionary nature, not used in practice? One reason, brought up in Opaluch et al. (2011), concerns the imperfect monitoring problem. The authors point out that in the past the government had disagreements with the firms about the size of their profits, and that, from the regulator's perspective, it is easier to account for production levels and revenues instead. Analysis of such conflicts is outside the scope of this paper. Another reason, however, may have to do with

revenues to the state or give up the business. Subsequently, the country had to design special incentives, such as minimum revenue guarantees, to attract foreign investment (Ghandi and Lin (2014)).

the dynamics of payoffs: revenue shares are collected by the state immediately when production begins, while profit shares would come only after firms recover their upfront costs of drilling. How land leasing can be arranged over time, and what types of contracts would be preferable for the government to use are interesting topics that I may address in the future work.

In general, investors are always interested in improving their revenues and reducing their costs. The strategies they pursue to achieve this goal often interact with the regulatory environment. While in this paper I consider bidding for leases and investment timing as such strategies, in other settings firms can engage in a variety of alternative activities that also influence their revenues and costs. In offshore leasing, the opportunity to learn from geographic neighbors has been shown to play a key role for investment incentives of the drillers (Hodgson (2018)). Onshore, the location of the well can be an important choice to make (Agerton (2019)). In the context of fracking technologies, innovation may become another dimension of the firms' strategies (Steck (2018)). How the design of leasing contracts may affect these complex firm behaviors is a possible future avenue for my research.

Although this paper focuses on the potential of revenue-sharing contracts to enrich the government and reduce investment of firms, there are a number of other features such contracts possess. One interesting opportunity these contracts provide has to do with collusion incentives. In oil and gas leasing in the US, firms are known to sometimes agree on rigging the bids in the public land auctions or on dividing the lands ahead of signing leasing agreements with private parties.[30] If, however, the landowner gains most of the benefit through the revenue sharing channel, the level of competition may prove to be relatively unimportant. Thus, collusion incentives become a part of the seller revenue — firm investment trade-off. Another feature of these contracts is reminiscent of the cost-plus agreements: the risk associated with the project is shared between the seller and the firm. The oil leasing environment, where the major source of uncertainty is the oil price, represents an attractive setting to study the importance of risk sharing for investment and revenue outcomes. The effect of lease contract design on collusive behavior and risk sharing is a promising direction for my future work.

[30] Western District of Oklahoma (2017), a filing in relation to the litigation against Chesapeake Energy, describes how two firms essentially divided among themselves millions of acres of oil and gas lands in Mississippi Lime Play, leasing them at the smallest possible royalty rates from the private landowners.

References

Agerton, M. (2019): "Learning Where to Drill: Drilling Decisions and Geological Quality in the Haynesville Shale," *Working Paper.*

Allcott, H., and D. Keniston (2017): "Dutch Disease or Agglomeration? The Local Economic Effects of Natural Resource Booms in Modern America," *The Review of Economic Studies*, 85(2), 695–731.

An, Y., and X. Tang (2019): "Specification and negotiation in incomplete contracts," *The RAND Journal of Economics*, 50(1), 57–92.

Anderson, S. T., R. Kellogg, and S. W. Salant (2018): "Hotelling under pressure," *Journal of Political Economy*, 126(3), 984–1026.

Arps, J. J. (1945): "Analysis of decline curves," *Transactions of the AIME*, 160(01), 228–247.

Asker, J., A. Collard-Wexler, and J. D. Loecker (2018): "Market Power, Production (Mis) Allocation, and OPEC," *American Economic Review (accepted).*

Athey, S., J. Levin, and E. Seira (2011): "Comparing Open and Sealed Bid Auctions: Theory and Evidence from Timber Auctions," *Quarterly Journal of Economics*, 126(1), 207–257.

Bajari, P., S. Houghton, and S. Tadelis (2014): "Bidding for Incomplete Contracts: An Empirical Analysis of Adaptation Costs," *American Economic Review*, 104(4), 1288–1319.

Bajari, P., and G. Lewis (2011): "Procurement contracting with time incentives: theory and evidence," *The Quarterly Journal of Economics*, 126(3), 1173–1211.

Bhattacharya, V., A. Ordin, and J. W. Roberts (2019): "Bidding and Drilling Under Uncertainty: An Empirical Analysis of Contingent Payment Auctions," *Working Paper.*

Board, S. (2007): "Selling Options," *Journal of Economic Theory*, 136(1), 324–240.

Brown, J., P. Maniloff, and D. Manning (2018): "Effects of State Taxation on Investment: Evidence from the Oil Industry," *Federal Reserve Bank of Kansas City Working Paper.*

Bureau of Land Management (2015): "Oil and Gas Leasing; Royalty on Production, Rental Payments, Minimum Acceptable Bids, Bonding Requirements, and Civil Penalty Assessments," *Federal Register.*

Che, Y.-K., and J. Kim (2010): "Bidding with Securities: Comment," *American Economic Review*, 100(4), 1929–35.

Cong, L. W. (2019): "Auctions of Real Options," *Working Paper.*

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

Djankov, S., et al. (2010): "The Effect of Corporate Taxes on Investment and Entrepreneurship," *American Economic Journal: Macroeconomics*, 2(3), 31–64.

Energy Information Administration (2016): "Trends in US Oil and Natural Gas Upstream Costs," *Report.*

Fanchi, J. R., and R. L. Christiansen (2017): *Introduction to petroleum engineering.* John Wiley, 1 edn.

Ghandi, A., and C.-Y. C. Lin (2014): "Oil and Gas Service Contracts Around the World: A Review," *Energy Strategy Reviews*, 3, 63–71.

Goldsworthy, B., and D. Zakharova (2010): "Evaluation of the oil fiscal regime in Russia and proposals for reform," *International Monetary Fund*, (10-33).

Goolsbee, A. (1998): "Investment Tax Incentives, Prices, and the Supply of Capital Goods," *The Quarterly Journal of Economics*, 113(1), 121–148.

Haile, P. A., K. Hendricks, and R. H. Porter (2010): "Recent U.S. offshore oil and gas lease bidding: A progress report," *International Journal of Industrial Organization*, 28(4), 390–396.

Haile, P. A., H. Hong, and M. Shum (2003): "Nonparametric tests for common values at first-price sealed-bid auctions," *National Bureau of Economic Research*, (w101105).

Hale, B. W., and R. K. Hall (2017): "Historical terminal decline rates review of unconventional reservoirs in the United States," *Society of Petroleum Engineers Unconventional Resources Conference.*

Hall, R. E., and D. W. Jorgenson (1967): "Tax Policy and Investment Behavior," *American Economic Review*, 57(3), 391–414.

Hausman, C., and R. Kellogg (2015): "Welfare and Distributional Implications of Shale Gas," *National Bureau of Economic Research*, (w21115).

Hendricks, K., J. Pinkse, and R. Porter (2003): "Empirical Implications of Equilibrium Bidding in First-Price, Symmetric, Common Value Auctions," *Review of Economic Studies*, 70(1), 115–145.

Hendricks, K., R. Porter, and G. Tan (1993): "Optimal Selling Strategies for Oil and Gas Leases with an Informed Buyer," *The American Economic Review*, 83(2), 234–239.

Hendricks, K., and R. H. Porter (1988): "An Empirical Study of an Auction with Asymmetric Information," *American Economic Review*, 78(5), 865–883.

——— (2014): "Auctioning Resource Rights," *Annual Review of Resource Economics*, 6(1), 175–190.

Herrnstadt, E., R. Kellogg, and E. Lewis (2019): "Royalties and Deadlines in Oil and Gas Leasing: Theory and Evidence," *Working Paper.*

Hodgson, C. (2018): "Information Externalities, Free Riding, and Optimal Exploration in the UK Oil Industry," *Working Paper.*

House, C. L., and M. D. Shapiro (2008): "Temporary Investment Tax Incentives: Theory with Evidence from Bonus Depreciation," *American Economic Review*, 98(3), 737–768.

Kellogg, R. (2014): "The Effect of Uncertainty on Investment: Evidence from Texas Oil Drilling," *American Economic Review*, 104(6), 1698–1734.

Kong, Y. (2017a): "Selective Entry: Estimation and Evidence," *Working Paper.*

——— (2017b): "Sequential Auctions with Synergy and Affiliation Across Auctions," *Working Paper*.

Kong, Y., I. Perrigne, and Q. Vuong (2019): "Multidimensional Auctions of Option Values," *Working Paper*.

Krasnokutskaya, E., and K. Seim (2011): "Bid Preference Programs and Participation in Highway Procurement Auctions," *American Economic Review*, 101(6), 2653–2686.

Lee, Y., and R. H. Gordon (2005): "Tax Structure and Economic Growth," *Journal of Public Economics*, 89(5-6), 1027–1043.

Male, F. (2019): "Using a segregated flow model to forecast production of oil, gas, and water in shale oil plays," *Journal of Petroleum Science and Engineering*, 180, 48–61.

Metcalf, G. (2018): "The Impact of Removing Tax Preferences for U.S. Oil and Natural Gas Production: Measuring Tax Subsidies by an Equivalent Price Impact Approach.," *Journal of the Association of Environmental and Resource Economists*, 5(1), 1–37.

Milgrom, P. R., and R. J. Weber (1982): "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50(5), 1089–1122.

Newell, R. G., B. C. Prest, and A. Vissing (2016): "Trophy hunting vs. manufacturing energy: The price-responsiveness of shale gas," *Technical report, NBER*.

Opaluch, J. J., et al. (2011): "Policies to Affect the Pace of Leasing and Revenues in the Gulf of Mexico," *Summary Report to the Department of the Interior*.

Osmundsen, P., K. E. Rosendahl, and T. Skjerpen (2015): "Understanding Rig Rate Formation in the Gulf of Mexico," *Energy Economics*, 49, 430–439.

Postali, F. A., and P. Piccheti (2006): "Geometric Brownian motion and structural breaks in oil prices: a quantitative analysis," *Energy Economics*, 28, 506–522.

Rabe, B. G., and R. L. Hampton (2015): "Taxing fracking: The politics of state severance taxes in the shale era," *Review of Policy Research*, 32(4), 389–412.

Rao, N. (2018): "Taxes and US Oil Production: Evidence from California and the Windfall Profit Tax," *American Economic Journal: Economic Policy*, 10(4), 268–301.

Reimer, M. W., M. Guettabia, and A.-L. Tanaka (2017): "Short-run impacts of a severance tax change: Evidence from Alaska," *Energy Policy*, 107, 448–458.

Robinson, M. S. (1984): "Oil Lease Auctions: Reconciling Economic Theory with Practice," *Discussion Paper #292*.

Steck, A. L. (2018): "Industry Dynamics with Social Learning: Evidence from Hydraulic Fracturing," *Working Paper*.

Western District of Oklahoma (2017): "Case No. CIV-16-209-HE, Document 164," *United States District Court*.

Yam, S., S. Yung, and W. Zhou (2012): "Optimal selling time in stock market over a finite time horizon," *Acta Mathematicae Applicatae Sinica*, 28(3), 557–570.

Zwick, E., and J. Mahon (2017): "Tax Policy and Heterogeneous Investment Behavior," *American Economic Review*, 107(1), 217–248.

Appendix

A. Sample construction details

The sample construction involves merging the datasets together, defining potential bidders, and selecting the leases likely intended for oil production but not gas production. The data comes from auction bid sheets, lease maps, well data obtained from the data vendor *Drillinginfo*, and oil and gas prices together with GDP deflator values from St. Louis Fed (Federal Reserve Economic Data). Bid sheets describe bid sizes, bidder names, lease name (if the auction is successful), tract name, assigned royalty rate, and some location characteristics for the leases. This dataset spans April 1994 through December 2015. The maps show exact location of each lease together with its name and span similar dates. Well data details location and basic characteristics of the wellbores: drill type (such as vertical or horizontal), production type (e.g. oil or gas), spud date, and several others. Well production is provided at wellbore level and describes output at up to 5 points in time: 6 months, 12 months, 24 months, 60 months, and up to las production date. For oil, WTI spot prices are used, and Henry Hub spots are used for gas. All nominal dollars (for bids and prices) and translated into 2009 USD using GDP deflator data, and the paper only uses these values throughout.

Valid lease names from the bid sheets are matched to the map data, resulting in 9432 matches. Auctions with misrecorded lease names are dropped. For successful matches, bidder names are cleaned using the OpenRefine software. The goal is to make sure names such as, for example, "Yates Petroleum Corp." and "YATES" are recognized as corresponding to the same bidder. This procedure results in 545 unique names, and identifiers corresponding to each of them are created. Next, wells are merged with the leases based on their locations.[31] Some lands in the data are resold. In those cases, mismatches between wells and leases are possible. To rule out such cases, I make sure that delays between sale date and well spud date are never negative, and they exceed the duration of 5 years only when there is at least one well matched to the lease drilled within 5 years from saledate.[32] Thus, some leases may have more than one well assigned to them.[33]

[31] For the wells, surface locations are used.

[32] There are several cases when delay from sale to spud for the earliest successful match exceeds 5 years by up to 2 months. Based on the conversations with a New Mexico State Land Office employee, the office used to be lenient with the firms, allowing the latter to sometimes slightly exceed the allotted duration of the lease. Thus, I regard such cases with delays slightly above 5 years as valid matches.

[33] Using map data for all leases, namely 12774 of them, I also check how many wells match to every lease. In total,

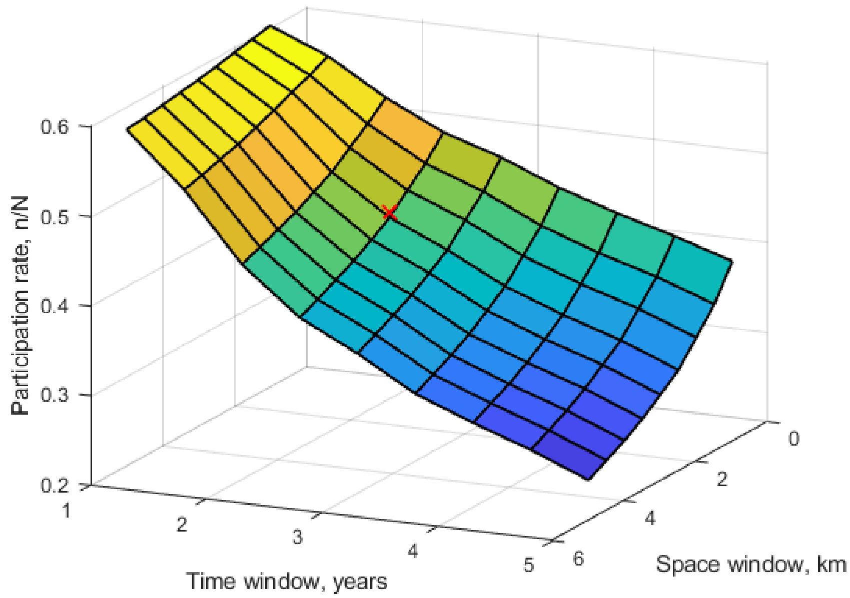

Figure 10: Actual-to-potential bidder ratio, or participation rate, n/N for various time and distance windows in definition of adjacent leases. The marker denotes the value corresponding to the definition used in the paper.

As the next step, I define potential bidders for each auction a. To do that, I consider auctions for leases within d kilometers that were held within t years of the auction a. For this collection of auctions, I group together all unique bidders. This group is made to always include one particular bidder, Yates Petroleum, a firm that has participation rate of roughly 50% across all auctions in the data.[34] The number of potential bidders N_a for a is defined as the number of unique bidders in this group. The space and time bands chosen in this paper are 2.5 years and 2.5 km (for comparison, most of these leases are for 0.5×1.0 miles plots of land). Figure 10 represents how average participation rate n/N changes across the full dataset depending on the choice of d and t.

Next, for each well I compute revenues from oil and gas. For each of the commodities, I determine the average price over the period of production and multiply cumulative output by that price. I then define the revenue ratio as the ratio of oil revenue to \$1 + gas revenue. Finally, each lease in the sample I mark as oil-rich if the median well within 5 km from the lease has revenue ratio

there are 1326 wells matched (for reference, there are about 38000 wells spudded since 1990 somewhere in New Mexico). Across the State Land Office leases, 1326 had at least one well matched to them. Of these, most leases have either exactly one well (880) or exactly two wells (233).

[34]The second most frequent bidder, Daniel E. Gonzales, has participation rate of just under 10%.

	(1) Bid 1	(2) Bid 1
P_0		1.144***
		(0.308)
P_g	2.262	-2.516
	(2.051)	(2.069)
N	682	682
R^2	0.332	0.354

Standard errors in parentheses
$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Table 6: OLS regressions for the winning bid with gas prices P_g and oil prices P_0. Gas prices available only from July 1997 on; consequently, 682 observations are used for the regression instead of the full sample. Regressions control for acreage, potential bidders, and year fixed effects.

exceeding the value of 5.[35] As a result, 3594 leases end up being marked as oil-rich. Looking at wells matched with such leases, the revenue ratio average is 18.1 with the median of 11.3. Less than a quarter of these wells have revenue ratio below 5. From the oil-rich lease sales, I select sealed bid auctions only, which roughly halves the sample.[36] Finally, I focuse on exploratory and discovery leases only, leaving out the development ones.[37] This leaves me with the final sample consisting of 786 observations. See Table 1 for the sample description.

To confirm that the firms also think of the selected leases as intended for drilling oil wells, I study the dependence between bids and gas prices (Henry Hub spot prices). Results are reported in Table 6. The coefficient on gas prices is insignificant for the winning bid even when not controlling for the price of oil. The magnitude of the coefficients also appears low compared to the magnitude of the coefficient on the oil price given the average gas prices of $4.5 for the corresponding time period. I also estimate these regressions with years 2005 and 2008 excluded (gas prices were exceptionally high during this period), but the result remains the same.

[35] Sometimes, the wells are viewed as oil-focused when the total energy equivalent of oil is at least 50% of the energy equivalent produced by the well — see, e.g., Anderson, Kellogg, and Salant (2018). My definition is based on revenues instead.

[36] The rest of the auctions are open outcry. The assignment of the auction format can be viewed as random — see Kong (2017a).

[37] Exploratory and discovery leases do not seem to differ from each other in anything except the royalty rate, which is 1/6 for the former and 1/8 for the latter. Development type is assigned by the state land office to leases for lands close to other already developed areas of the state and whose chances to contain oil are well understood. By contrast, exploratory and discovery types are assigned to leases at the geographic frontier of land sales and whose adjacent areas are less developed.

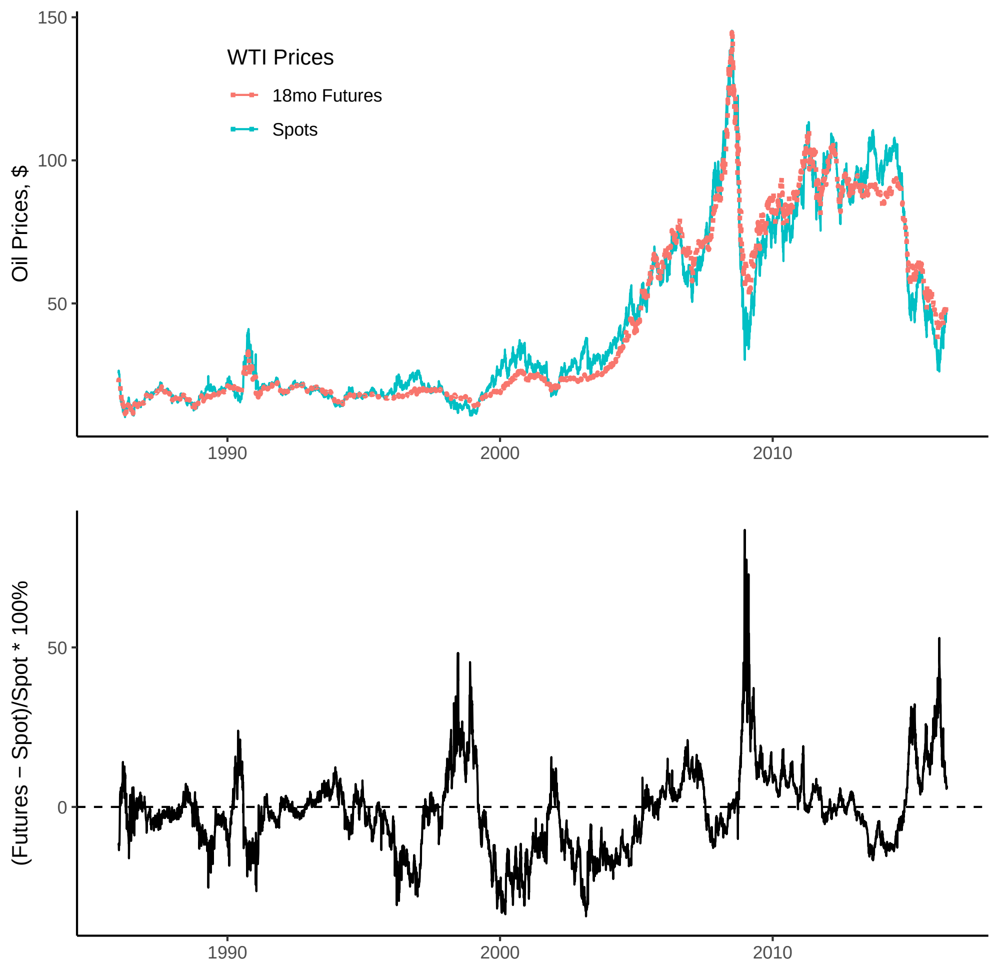

Figure 11: Differences between spots and futures prices: nominal (top) and percentage (bottom).

B. Oil Prices

B.1. Spots vs Futures

The market for upstream oil is notoriously opaque. A firm may use a variety of strategies to sell oil, relying both on spot sales and long-term contracts. Whether it is spot prices or futures for oil that best explain the behavior of the firms to an extent could be addressed within the framework considered in this paper.

To start, note the differences between the two time series of WTI spots and WTI futures with maturity of 18 months presented in Figure 11.[38] Although by and large the two processes are very close to each other, there are time periods when they diverge rather far apart. This variation allows

[38]Price data is from Bloomberg.

	(1) Bid 1	(2) Bid 1	(3) Log(Bid 1)	(4) $\mathbb{1}$(drill)	(5) $\mathbb{1}$(drill)
P_0		-0.250		-0.188	-0.022
		(0.943)		(0.403)	(0.458)
P_f	1.491***	1.817		0.537	0.375
	(0.443)	(1.473)		(0.590)	(0.676)
Log(P_0)			0.955**		
			(0.406)		
Log(P_f)			-0.553		
			(0.628)		
N	786	786	786	786	506
R^2	0.374	0.375	0.483	0.135	0.169

Standard errors in parentheses

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Table 7: OLS regressions with oil prices P_0 and 18 month oil futures P_f. Bids are in \$1,000. $\mathbb{1}$(drill) is measured in %, taking either 0 or 100 as values. Column (5) controls for the runner-up bid. All regressions include controls for year fixed effects, acreage, and the number of potential bidders.

to compare which of the prices better explains the drilling observed for the leases in the sample.

Table 7 shows some basic regressions for bids and drilling decisions. Column 1 provides the result for a basic regression similar to that of Table 2 where WTI futures are used instead of spots. The coefficient for the futures is statistically significant and economically meaningful: a \$1 increase in the prices is associated with a \$1,491 rise in the winning bid. When controlling for both contemporaneous oil prices and the futures prices (column 2), the coefficients for both lose their significance. The magnitude of effect, however, is strong: for instance, even if both prices go up by \$1 just before the auction, the winning bid would on average increase by \$1,567. The loss of significance is symptomatic of the collinearity issue: as mentioned above, the two prices are clearly strongly related. In the log-version of this regression (column 3), however, the coefficient on P_0 retains significance after inclusion of the futures. The table also shows the relationship between the drilling decisions and the prices (columns 4 and 5), with conclusions similar to the bidding regression. As discussed throughout the paper, however, the drilling decision is contingent on the entire realization of prices after the auction, rather than just prices at the auction. To check which of the price paths is more consistent with the observed drilling decisions, I further compare the prices to each other in the exercise in Appendix D.

	Spots	Futures
μ	0.0355	0.0433
	(0.0663)	(0.0226)
σ	0.3038	0.1990
	(0.0131)	(0.0076)
N	360	360

Table 8: Geometric Brownian Motion estimation. Price data starts from January 1986.

B.2. Price Parameter Estimation

In this subsection, I provide the estimates for the Geometric Brownian Motion based on spots and futures prices. The estimation is straightforward: the likelihood is written down based on equation

$$\log\left[P_{t+\Delta}\right] - \log\left[P_t\right] = (\mu - \sigma^2/2)\Delta + \sqrt{\Delta}\varepsilon_t \tag{20}$$

where ε_t are i.i.d. standard normal innovations and $\Delta = 1/12$ corresponds to one month (i.e. t is measured in years). The results are provided in Table 8. Compared to the spots, oil futures have lower volatility; for the optimal stopping problem, this results in overall lower threshold prices and thus earlier drilling.

C. Well production extrapolation

The data for wells comes from *Drillinginfo* and includes location and production observations for wells drilled between January 1990 and January 2019. Specifically, production is available for each wellbore[39] at up to five different points in time: 6 months, 12 months, 24 months, 60 months from production start, and cumulative output to date. The data covers both oil and gas output. In total, there are 38644 wells in the set corresponding to 39923 wellbores. Of these, 2443 wells match to 1326 unique state leases.

In the paper, this full dataset, as opposed to exclusively the wells matched to sample leases, is used for a variety of purposes. I look at the production rates of the wells to see if they respond to prices. This is important since one of the main assumptions of the model is that the agents do not control the output curves of the wells, only the timing of drilling. I discuss how informative the

[39] Wellbores drilled in different directions starting from the same spot on the surface together constitute a well.

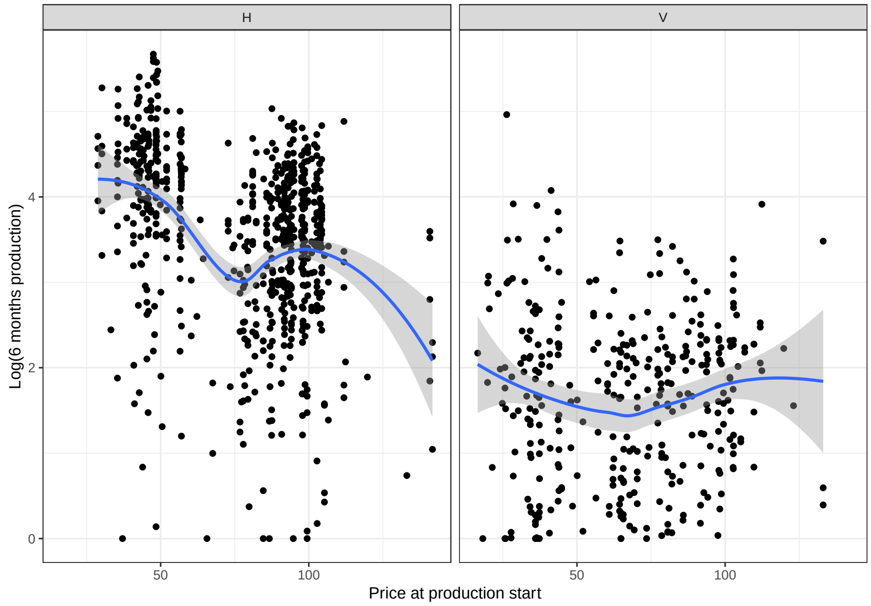

Figure 12: Well output after 6 months compared to the price at production start. Grouped by horizontal and vertical wells. Only wells corresponding to state leases are used.

observations about early production of the wells are about their lifetime outputs. Finally, I set up a simple model predicting long-run production of the wells.

C.1. Production rates and prices

If oil producers could influence the rates of production and higher prices motivated them to accelerate the oil extraction, one would expect that output rates in the data are correlated with prices. As a simple check, I look at the output levels after 6 months of production and compare them to various measures of oil prices. Figure 12 illustrates one such exercise where the incentive measure is prices at the beginning of the well's production. For this and similar measures, such as, for example, average prices over the first 6 months, I observe that no pronounced upward trend exists in the data. Similar results can be obtained when looking at output after 12 months. Of course, this does not rule out the possibility that producers control decay curves of the wells completely, since this exercise does not address the selection problem (higher prices may trigger drilling of unproductive wells), and the measures of prices are somewhat noisy.

	Nobs	Mean	StDev	25%	50%	75%
Q_6/Q	17756	0.174	0.171	0.067	0.131	0.219
Q_{12}/Q	17756	0.276	0.215	0.135	0.228	0.355
Q_{24}/Q	17266	0.399	0.220	0.260	0.373	0.514
Q_{60}/Q	17266	0.629	0.228	0.501	0.636	0.788

Table 9: Early output share in cumulative output. Q_t is the oil output after t months. Only for wells drilled between 1990 and 2010.

C.2. Output prediction

As the first exercise for output prediction, I check how much of total well output corresponds to early production data. To this end, I restrict myself to wells that have been drilled prior to January 2010; this way, the observations describing cumulative output to date are likely to be close to the long-run level. Table 9 reports the corresponding statistics. It can be seen that after two to five years of production the output data becomes a farily reliable predictor of lifetime output, the latter being about 60% larger than the 60 months total. On the other hand, there is significant spread in early production data. For example, the interquartile range for the first 6 months of production relative to the cumulative level is 0.067 to 0.219.

To set up a model for predicting oil output, I reference the literature on decay curve analysis which traditionally has been applied to predict productivity of wells. The standard reference here is Arps (1945), where decay rate of output is parameterized as $D_t'/D_t = \alpha D_t^{\beta}$. Different values of the power parameter $\beta = 0, \beta \in (0, 1)$, and $\beta = 1$ give rise to exponential, parabolic, and harmonic models respectively. All three are commonly used in the modern analysis of well output (Fanchi and Christiansen (2017), Hale and Hall (2017)). These are the models I estimate.[40]

For either of the three decay rate parameterizations well output at time t can be represented as $Q_t = Qf(t)$ with $f(\infty) = 1$. Below, estimated parameters include well-specific Q and decay functions $f(t)$ which are assumed to be common across the wells.

1. Exponential model: $Q_t = Q(1 - \exp[-\beta t])$
2. Semi-exponential model: $Q_t = Q\Gamma_{inc}(\beta t, \alpha)$ where $\Gamma_{inc}(x, k) = \int_0^x z^{k-1} \exp(-z)dz/\Gamma(k)$

[40]One alternative production model in recent consideration characterizes wells as having multiple stages of operation (Male (2019)). Considering that with this approach the last stage of operation is typically assumed to be exponential decay and that this approach requires a rich dataset for well-level output, for my purposes I rely on the simpler models.

Model	Mean	StDev	Min	P25	P50	P75	Max	Fit	E[Q]	Median[Q]
Vertical 1	20.78	41.38	0.42	2.78	7.78	21.39	1122.79	0.72	43.68	24.41
Horizontal 1	82.77	84.74	0.54	23.56	56.86	113.88	681.42	0.95	158.37	123.24
Vertical 2	20.14	40.99	0.46	2.76	7.97	20.48	1469.80	0.74	11365.93	6414.31
Horizontal 2	80.51	78.85	0.55	24.34	58.09	111.40	694.48	0.96	220.43	169.99
Vertical 3	20.24	40.46	0.45	2.77	7.92	20.83	1283.84	0.73	19.7×1E6	11.1×1E6
Horizontal 3	80.81	79.63	0.52	24.21	57.90	111.86	701.31	0.96	21.3×1E6	16.5×1E6
Vertical data	22.49	56.52	0.00	1.15	6.31	20.47	2516.15			
Horizontal data	81.30	83.06	0.00	22.88	57.79	113.24	886.22			

Table 10: Production model predictions. Reference data is for cumulative oil to date. Fit is the share of variation in Q_t that is explained by the model prediction $\hat{Q}_t$ (note that the estimation is based on $\log(Q_t)$). The last two columns are forecasted lifetime quantities. Sample size: 19944 vertical wells, 5066 horizontal wells.

3. Parabolic model: $Q_t = Q(1 - (1 + \alpha t)^{-\beta})$

Denoting $Q_{i,t}$ the output at time t of well i, each model is estimated via

$$\min_{\theta} \sum_i \sum_t \Big(\log(Q_{i,t}) - \log(\hat{Q}_{i,t})\Big)^2 \tag{21}$$

This way, I avoid placing too much emphasis on the last observation corresponding to cumulative oil. Results are reported in Table 10. The exponential model is the only model with sensible predictions for the values of lifetime Q. In the paper, I use it for the derivations based on predictions of Q. The parameter estimates for this model are $\hat{\beta} = 0.03$ for vertical and $\hat{\beta} = 0.052$ for horizontal wells. To interpret, this implies that the rate of production slows down each month by 3% and by 5% for vertical and horizontal wells respectively.

D. Optimal stopping in the data

The characterization of the optimal stopping solution as described in Section 3.2 has an important implication: the data on drilling delays and prices can be used to test whether observations for the drilled wells are consistent with the model. Given the monotonicity of stopping paths in per-unit

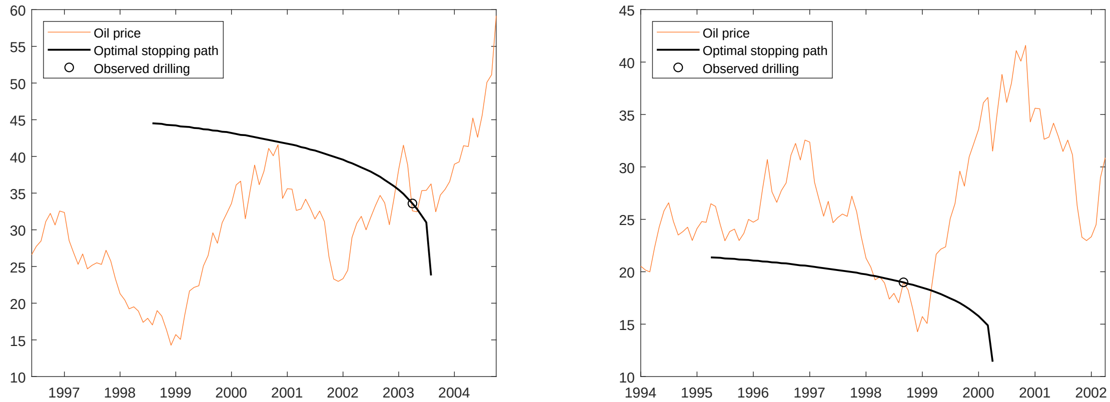

Figure 13: Model violations: mild (left) and strict (right). The model assumes that drilling happens the moment prices reach at least as high as the stopping path.

costs of drilling, there is exactly one path corresponding to each spudded well observation. This path would be inconsistent with the model whenever it predicts that drilling should have happened earlier than it did in the sample. Consider, for instance, Figure 13. The plot on the right provides an example where, according to the model, drilling should have been immediate in the first half of 1995. In reality, it was delayed by over 3 years, even as the oil prices remained consistently above the threshold path. In the same figure, on the left, the timing of another well also diverges from the model. Specifically, the stopping path suggests that drilling should have happened a few months earlier than it actually did.

The two examples suggest that there are several ways to evaluate the discrepancy between the model and the data. First, check every single case of drilling for a violation. Alternatively, focus on the more dramatic violations. To do the latter, I consider two classes of incompatibilities between the data and the model. A mild violation is when there have been at least three months prior to observed drilling when the real prices were above the stopping path. A strict violation is when there have been at least three[41] *consecutive* months prior to observed drilling when the real prices were above the stopping path. Additionally, for each case of a violation other statistics can be computed, such as the length of time between observed drilling and the first month of prices above the stopping path, or the overall number of months when the drilling predicted by the model does not happen.

The summary for violations is reported in Table 11. For spot prices, in about one third of the

[41]The amount of time it takes to obtain a drilling permit in New Mexico for an SLO lease historically has been about 2-3 months.

Statistic	Value (spots)	Value (futures)
Num wells	78	78
All violations	40	36
Mild violations	26	21
Strict violations	25	16
Violation mean length	23.5	23.6
Mean num violation months	7.2	6.3

Table 11: Violations summary. One version is based on spot prices P_0, the other — on 18 month futures prices P_f. The last two statistics are across all violations.

cases the model predictions are significantly different from what happens in the data. Most of these occur after July 2008, the date when the oil prices reached their historical high and then fell off by about 70%. When futures prices are used instead (see Appendix B for futures-spot comparison), the number of strict violations is significantly reduced. Furthermore, relative to spots, the overall extent to which the violations are incompatible with the model is lower: there are fewer months with observed prices above the implied stopping threshold.

Each violation is essentially the case of lease holders waiting longer than the model predicts they should have. Such behavior, however, can potentially be explained by other considerations. One issue with the modeling approach is that a simple Geometric Brownian Motion is used to predict the future prices. If, for instance, decisions of agents are based in part on implied volatility of the prices, it could possibly be understood why agents were waiting for higher prices: in my model, volatility is just a constant, whereas in reality oil price volatility increased post 2008. Another consideration has to do with capacity constraints. The opportunity to drill can be exercised in real life only when there is rig available for drilling. In the model, rig shortages can be interpreted as high costs; the caveat, however, is that in the model costs are persistent and only shift due to price changes.