

Royalties and Deadlines in Oil and Gas Leasing: Theory and Evidence

Evan Herrnstadt
Ryan Kellogg
Eric Lewis*

February 21, 2019

DRAFT: COMMENTS WELCOME

Abstract

This paper seeks to estimate the impacts and explain the presence of two pervasive features of oil and gas lease contracts between mineral owners and extraction firms: the royalty and the primary term. The royalty is a percentage of hydrocarbon revenue that is paid to the mineral owner, and the primary term specifies the maximum number of years within which the firm must drill and produce from at least one well, lest it lose the lease. Using detailed data on lease contracts and the timing of drilling, we show empirically that primary term expiration dates have an economically significant impact on firms' drilling decisions: a large share of wells are drilled just prior to expiration. We then develop a model to explain why primary terms and royalties can help maximize the mineral owner's expected revenue from a lease, despite the distortions they generate. The royalty helps the mineral owner extract some of the information rents of the firm but also delays drilling; the primary term partially mitigates this moral hazard problem by encouraging earlier drilling. We examine how these contracts affect drilling and the payouts to mineral owners and firms in the Haynesville Shale formation of Louisiana.

*Herrnstadt: Compass Lexecon; eherrnstadt@gmail.com. Kellogg: Harris Public Policy, University of Chicago; and National Bureau of Economic Research; kelloggr@uchicago.edu. Lewis: Economic Analysis Group, Antitrust Division, United States Department of Justice; erikylewis@gmail.com. We thank William Patterson, Pengyu Ren, Grant Strickler, and especially Nadia Lucas for outstanding research assistance. We also thank conference and seminar participants at Aalto, AERE, Arizona, Columbia, DOJ, Federal Reserve Board, Harvard, Mannheim, Michigan, Northwestern, Penn, Penn State, the Southern Economic Association, Toulouse, Washington University, and the Winter Marketing-Economics Summit for valuable comments. The analysis and conclusions expressed herein are solely those of the authors and do not represent the views of the United States Department of Justice or Compass Lexecon.

1 Introduction

Because the owners of subterranean oil and gas often lack the expertise or financial capital necessary to extract their resources, they typically write contracts with specialized extraction firms to act as their agents. In the United States, as well as several other countries, these contracts take the form of mineral leases that ubiquitously contain royalty and “primary term” clauses that have the potential to distort the extraction firms’ incentives to drill and complete oil and gas wells. This paper is aimed at understanding both the economic rationalization for these clauses and their effects on firms’ activities, with an emphasis on the dramatic recent expansion in U.S. oil and gas production from shale resources.

An oil and gas lease grants an oil and gas firm an option, but not the obligation, to drill wells on the mineral owner’s parcel of land and extract the hydrocarbons. Upon signing a lease, the firm will pay the mineral owner a flat fee “bonus” that in some areas has exceeded \$10,000 per acre (Bogan 2009). The primary term specified in the lease contract is the period of time (typically 3 to 10 years) the firm has to drill at least one well on the leased parcel. If the firm does not commence production by the end of this primary term, the lease terminates, and the mineral owner is then free to sign a new contract with another firm or recontract with the original firm, likely requiring another bonus payment. If, however, the firm begins production of oil and gas during the primary term, the lease is “held by production” and enters a secondary term that lasts until the firm ceases production. During the secondary term, the firm may also drill additional wells on the parcel to increase its overall production rate.¹ Finally, the royalty specified in the lease dictates the percentage of the lease’s oil and gas revenue that the firm must pay to the mineral owner. These royalties are often significant, as the royalty rate typically lies between 12.5% and 25%. Brown et al. (2016) estimate that royalty payments associated with the six largest U.S. shale plays totaled \$39 billion in 2014.

The royalty and primary term clauses clearly distort firms’ incentives regarding how


¹We discuss the pooling of leases together into units in Section 2.

many wells to drill, when to drill them, and how much effort to invest in their fracking and completion. The incentive to drill at least one well to commence production prior to the expiration of the primary term has received considerable attention within the industry. Following the drops in natural gas prices in 2009 and in oil prices in 2014, there have been numerous reports of firms drilling unprofitable wells for the purpose of holding their lease acreage. For instance, the *San Antonio Express News* reported in 2012 that “Bobby Tudor, chairman and CEO of Tudor, Pickering, Holt & Co. LLC, an energy-focused investment bank in Houston, said many companies . . . are drilling quickly simply to meet the terms of their contract and keep their leases—not because they want to drill gas wells now.” Although royalties are less prominent in the news, they also distort firm decisions. The royalty is a tax on revenue only, thus driving a wedge between the firm’s profit and social surplus.

In this paper, we begin by providing empirical evidence in support of anecdotes like the one above. Our analysis focuses on the Haynesville Shale—an onshore natural gas-producing formation in northwest Louisiana that we describe in section 2. We collect detailed data on leases, drilling activity, fracking intensity and natural gas production. Figure 1 shows time series of natural gas prices, leasing, and drilling activity in the Haynesville region. The graph shows that the natural gas price and leasing peaked in early 2008, but that drilling did not peak until about two years later, shortly before many leases were to expire. This pattern of drilling suggests that an economically meaningful share of Haynesville drilling activity was driven by firms seeking to hold their leases, and motivates us to take a closer look at the data.

In section 3, we show that a substantial share of Haynesville wells were drilled just prior to lease expiration. In particular, there is a clear discontinuity in the probability of drilling a well at the lease expiration date. We show that the large number of parcels in which a single well was drilled just prior to expiration is difficult to fully explain with other factors such as cross-parcel information or common pool externalities, such as those discussed in Hendricks and Kovenock (1989), Hendricks and Porter (1996), and Lin (2013).

Figure 1: Time series graph examining the Henry Hub, Louisiana natural gas spot price, the number of leases signed in Haynesville parishes, and the number of Haynesville wells over time



Note: Source of Henry Hub natural gas prices is the Energy Information Administration. The source of the lease counts is DrillingInfo. We include all leases signed in the following parishes: Bienville, Bossier, Caddo, DeSoto, Natchitoches, Red River, Sabine, and Webster. The source of Haynesville well spuds is the Louisiana Department of Natural Resources (DNR).

Because the bunching of drilling at primary term expiration is clearly distortionary relative to a socially optimal drilling program, we next explore theoretically why mineral owners would include these contractual terms in their leases. We begin in section 4 by developing an analytic mechanism design model to illustrate the tradeoffs that a mineral rights owner and a firm face, building off of insights from Laffont and Tirole (1986) and Board (2007). In our model the firm has a hidden signal of the productivity of the lease, which leads to information rents. The firm, if it signs a lease, chooses when to drill and can also exert hidden costly effort (e.g., the quantity of fracking inputs such as sand and water) that determines how much is produced. Ex-post oil and gas revenues are contractible, but drilling and completion costs are not, so that the contract can be contingent on revenues and production but not on costs nor profits. Prices evolve stochastically.

We solve for the contract that maximizes the mineral owner's payoff, subject to the firm's incentive compatibility constraint and incentive rationality constraints. The mineral owner offers the firm a menu of contracts that include three types of payments: a signing bonus paid by the firm to the owner at the start of the contract, a drilling subsidy that is paid by the owner to the firm at the time of drilling, and a royalty payment on ex-post revenues paid by the firm to the owner. Per intuition from Riley (1988), Hendricks et al. (1993), and Bhattacharya et al. (2018), the royalty serves to reduce the firm's information rents but also delays drilling and reduces the firm's hidden effort. Our contribution is to show that, because the date of drilling is contractible, the owner can partially mitigate the royalty-induced drilling delay by paying a drilling subsidy to the firm at the time of drilling.

Our analytic model provides a rationale for why mineral owners include royalty clauses in their leases and for why, given the royalties, they also have an incentive to include clauses that induce the firm to drill sooner. A divergence between our analytic model and the contracts we observe in practice, however, is that we observe primary terms rather than drilling subsidies. This divergence can be rationalized by the existence of liquidity constraints on behalf of mineral owners, especially given that, as we show later, optimal drilling subsidies

are millions of dollars per well. But are primary terms then a viable second-best tool for mineral owners to increase their expected revenues, relative to a royalty-only infinite-horizon lease?

We address this question using a computable model in which the mineral owner makes a take-it-or-leave-it offer to the firm that includes a bonus payment, royalty, and possibly a primary term. We present this model in section 5. As with the analytic model, the firm has private information on expected gas production. We model the firm's drilling timing decision as an optimal stopping problem given the contract it signs and a common knowledge stochastic process for both gas prices and drilling costs. Our modeling of this problem builds off of prior work on the determinants of drilling timing, including Kellogg (2014), Agerton (2018), and Bhattacharya et al. (2018). An important new feature of our model is its incorporation of the possibility that the firm and mineral owner may agree to extend the lease upon expiration of the primary term (i.e., in the event that drilling does not occur), subject to payment of a second bonus that is contingent on the state variables (gas price and drilling cost) at the expiration date.

In section 6, we calibrate our computational model using drilling and production data from the Haynesville Shale and historical natural gas prices and drilling costs. We then use the model in section 7 to explore how varying the royalty and primary term affect both the mineral owner's and firm's expected values, as well as drilling outcomes, beginning with a simple case in which the lease can only accommodate a single well. We find that the optimal contract includes not just a royalty but also a finite primary term, and that including the primary term improves the mineral owner's expected revenues by 4% (\$90,000) relative to a royalty-only contract. When we examine drilling programs, we see that the primary term indeed pulls expected drilling forward in time, but at the cost of creating a discontinuous fall in the drilling probability at the expiration date. Such a drop is of course not observed in the socially optimal drilling program, in which the drilling probability falls gradually over time.

If we relax the mineral owner’s liquidity constraint and allow it to offer a contract that includes a state-invariant drilling subsidy, but no primary term, we find that the mineral owner is better off, consistent with our analytic model. The optimal drilling subsidy increases the owner’s expected revenues by 10% (\$225,000) relative to the royalty-only contract, and it delivers a time path of drilling probabilities that declines smoothly over time and therefore better approximates the socially optimal path (though there is still a delay on average, consistent with the analytic model).

Finally, we enrich the model to examine the effects of primary terms when, as is common in the shale era, a single well can hold a lease (or collection of leases) upon which additional future wells may be drilled. We show that primary terms in this situation no longer provide a substantial revenue benefit to the mineral owner, for two reasons. First, the drilling timing distortion induced by the primary term is large because the firm almost always prefers to drill rather than pay a large bonus to extend the lease (since the bonus will account for the option value of drilling all of the wells, not just one). Second, the primary term does not provide any incentive to accelerate drilling for any well other than the first one, so that later wells remain fully exposed to the royalty’s delay incentive. In contrast, we show that a drilling subsidy still substantially increases the mineral owner’s expected payoff. These results suggest that leasing mechanisms that were consistent with revenue maximization for historic one-well leases have not kept up with the changing economic and institutional environment instigated by the development of shale oil and gas.

Our paper is predated by an extensive literature examining auction design for goods with uncertain value (see Porter (1995) and Haile et al. (2010) for summaries in oil and gas). Our paper builds off of the insights of Riley (1988), Hendricks et al. (1993), and Bhattacharya et al. (2018), who point out that a non-zero royalty increases expected revenue to the mineral owner when firms have private information about the deposit’s underlying productivity. Bhattacharya et al. (2018) is closest to our work, in that it uses a computational model of firms’ drilling timing problem to evaluate optimal oil and gas royalties for state-owned

parcels in New Mexico, focusing on the tradeoff between reducing firms' information rents and reducing moral hazard.

Our paper is distinct in that it explicitly models the effects of primary term length on drilling decisions and on the value received by the mineral owner and firm, accounting for the possibility that the owner and firm may re-contract (for an additional bonus payment) following lease expiration. We thereby provide what we believe is the first rationalization for the inclusion of both royalties and primary terms in oil and gas leases, and we compare these contracts to other possible mechanisms, such as direct subsidization of drilling. Our findings may shed light on similar contracts in other settings. For example, franchisee contracts often include royalty payments to the franchisor and impose a finite period of time for the franchisee to open a minimum number of franchise units (Kalhins 2005). Our model also links to the broader literature that considers asset sales involving contingent payments, recently surveyed in Skrzypacz (2013). For example, Cong (2018) considers a model in which the decision of when to sell the option itself is endogenous, and finds that this too is delayed relative to the social optimum.

2 Institutional background: Louisiana and the Haynesville formation

2.1 The Haynesville Shale

Our focus is on the Haynesville Shale formation, a tight shale gas formation in northwest Louisiana and east Texas. The development of new techniques combining horizontal drilling with hydraulic fracturing made it profitable to extract from the Haynesville formation, and speculation and drilling in the Haynesville exploded in early 2008. The same technology led to drilling booms in other shale formations throughout the United States, including the Bakken of North Dakota (oil), the Marcellus in Pennsylvania (gas), and the Barnett (gas)

and Eagle Ford (both oil and gas) in Texas.

We focus our study on the Haynesville Shale—and specifically the Louisiana portion of the Haynesville—for two reasons. First, the Haynesville produces almost exclusively dry natural gas. The near-absence of natural gas liquids and crude oil allows us to focus our analysis on a single output. Second, the economic and legal institutions in Louisiana that shape mineral leasing and the pooling of leases into units facilitate our empirical work, which requires us to match wells to their pooling units and associated leases. We summarize these institutions below and provide additional detail in appendix A.

2.2 Institutional details

When a firm is interested in drilling on privately-owned land, it must negotiate a lease with the mineral owner.² In almost all cases, this lease has a three-part structure. First, the firm pays the mineral owner a cash bonus when the lease is signed. Second, the lease specifies a royalty rate, which is the fraction of the revenue of a well that is paid to the mineral owner. Third, the lease specifies a primary term, which is a set amount of time that the firm has an option to drill and commence production before it loses the lease. If the firm drills a productive well before the lease expires, the lease is “held by production”, which means that the firm continues to hold the lease as long as there is commercial oil and gas production on the lease. Leases may also include extension clauses, which specify that the firm can pay a set amount of money to extend the primary term for a set amount of time.

In practice, leases typically have a continuous operations clause that allows the firm to spud—i.e., start, but not necessarily complete—a well prior to expiration and still hold the lease as long as operations are “continuous” as defined in the lease contract. Once continuous operations cease, the well must produce in order to hold the lease; thus, the firm will eventually have to complete the well to hold its leased acreage. While there have

²The lease structure we describe here also applies nearly ubiquitously to publicly-owned oil and gas. The primary difference between private and public leasing is that public leases are usually allocated by well-organized bonus bid auctions (rather than unstructured negotiations), in which the royalty and primary term are fixed in advance, either by statute or regulation.

been some reports that even more preliminary steps such as building a road or securing a drilling permit are sufficient to activate a continuous operations clause, case law and our discussions with industry experts suggest that such instances are infrequent in the Louisiana Haynesville. Thus, we will focus on spudding a well as the necessary step to hold a lease beyond its primary term.

One problem for firms and regulators is that leases are typically small relative to the area that is drained by a single well. This problem is especially severe for horizontal wells, which may have horizontal bores of 5,000 feet or more. Therefore, state regulators have established laws on how leases are to be combined into pooling units. In Louisiana, the default pooling unit for the Haynesville Shale formation is the square-mile section from the Public Land Survey System (PLSS). Because horizontal wells in tight shale formations primarily recover natural gas that is locked within rock close to the well bore, square-mile units have space for multiple horizontal wells that run parallel to one another, each with a length that is just shy of one mile (5,280 feet).

Drilling a Haynesville well within a Haynesville pooling unit holds all current leases within the pooling unit. This means that the operating firm has the right to drill any additional Haynesville wells within the same unit. When there are multiple firms holding leases within a given pooling unit, drilling operations effectively function as a joint venture. One lead firm, typically the one with the highest acreage share of leases, becomes the operating firm and decision maker. Costs and revenues are distributed to all leaseholding firms on an acreage-weighted basis. Each firm then has the obligation to distribute royalties on revenues to each of the relevant mineral owners on an acreage-weighted basis.

Owners of mineral rights that are unleased at the time of drilling—either because their parcels were never leased or because their leases expired prior to drilling—effectively become participants in the joint venture with acreage-weighted shares in the profits. Because mineral owners typically do not have the financial liquidity to pay their share of the drilling and completion costs, Louisiana statute (LA R.S. 30:10) provides them the option not to pay. In

that case, they do not receive their share of revenues until the well’s overall revenues cover its costs (i.e., the well “pays out”). Firms, then, cannot earn strictly positive profits from unleased acreage (except to the extent that they can “pad” costs in their cost reports to the unleased mineral interests), even as they remain exposed to the risk of strictly negative profits if the well does not pay out.

Thus, it is the conversion of acreage within a unit from leased to unleased mineral interest that provides firms with the incentive to drill prior to the expiration of primary terms. Any given unit will typically consist of many leases, not all of which expire at the same time. The drilling incentive provided by a given lease’s pending expiration then depends on the acreage of that particular lease and the upcoming schedule with which the remaining leases will expire.³

3 Data and descriptive empirics

In this section we first discuss our data sources, sample selection, and summary statistics. We then present descriptive evidence on primary terms that motivates our theoretical model.

3.1 Data sample and summary statistics

Our analysis uses data on leasing, wells, and production. From DrillingInfo, an oil and gas industry intelligence service, we compile data on the universe of oil and gas leases in Louisiana that started between 2002 and 2015. These lease data include the start date of the lease, the primary term, any extension options, the royalty rate, the sections of the lease,⁴ and the acreage of the lease. Well data are taken from the Louisiana Department of Natural Resources (DNR) and include information on permit dates, spud dates, completion

³For instance, if the schedule of lease expirations is such that only one small lease expires today, with all other acreage expiring two years from now, the expiration of the small lease today provides only a small incentive to drill immediately.

⁴Leases that span multiple sections do not list acreage for each individual section. We therefore impute the section-specific area of the lease for such leases.

dates, the volume of water used in hydraulic fracturing, and whether the well targets the Haynesville formation. We obtain well-level production data from DrillingInfo. A detailed discussion of these data sources, our data cleaning procedures (especially for the lease data), and our matching of wells to sections and to leases, is provided in appendix B.

Because pooling unit boundaries in the Louisiana portion of the Haynesville Shale are typically PLSS sections, we use these sections as the unit of observation.⁵ We limit our sample to sections that became Haynesville units. In addition, we drop sections where there appears to have been confounding activity that may affect leasing or drilling incentives, such as leasing that preceded the Haynesville boom or that had non-Haynesville production.⁶ There are 1,197 sections in our sample which are mapped in panel (A) of figure 2.

Table 1 shows summary statistics for these sections. Sections tend to have their first lease expire between 2008 and 2012, with a median of 2009. Seven hundred thirty-three sections (61%) have Haynesville wells drilled. Of the sections with drilling, 72% had only one well drilled, 15% had 2 wells drilled, 5% had 3 wells drilled, and 4% had 4 or more drilled. The most wells we observe in a single section is 19. The initial Haynesville well drilled in a section tends to be spudded between 2008 and 2011. Panel (B) of figure 2 shows the number of wells drilled per section.

Table 2 shows descriptive statistics about leases within that sample. We observe 34,904 leases in our section sample, and the average section has about 29 leases. Leases typically started between 2005 and 2011. Most leases have 36 month primary terms, and royalties are typically between 19% and 25%, with 25% being a very common royalty rate. About 78%

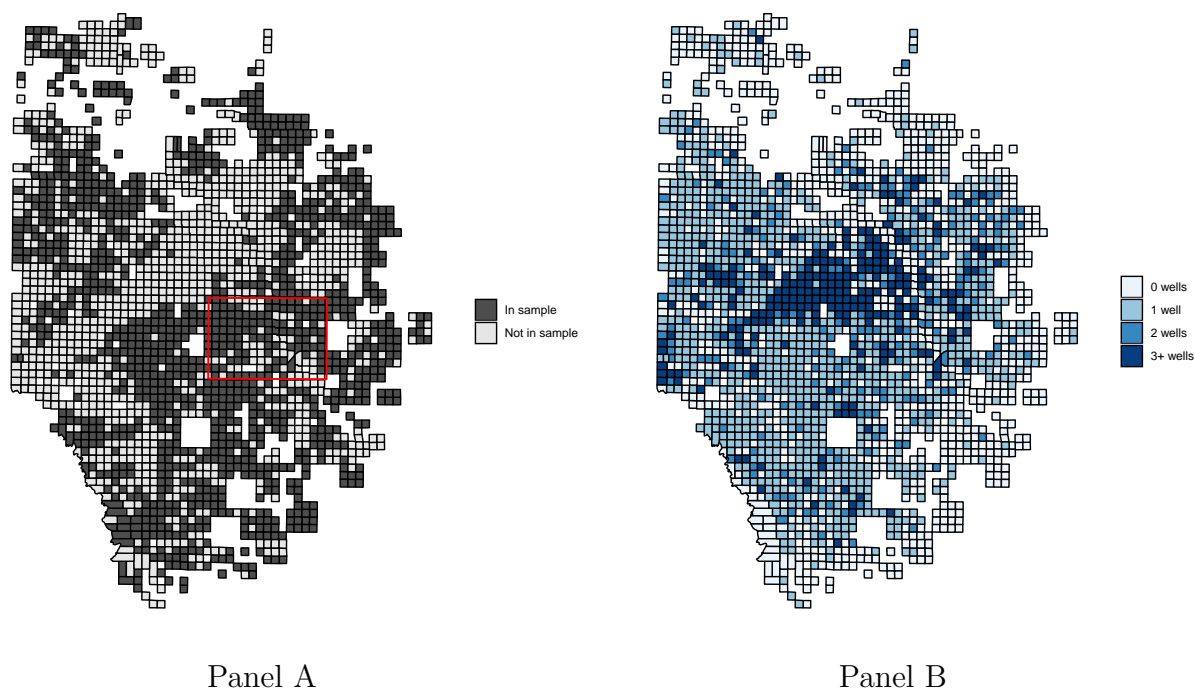
⁵We use “section” to refer to PLSS sections throughout the paper.

⁶Specifically, we do the following: First, we limit our analysis to sections that became Haynesville units, ensuring that our analysis is limited to the area where the Haynesville Shale was known to be potentially productive. Second, we limit our analysis to sections in which leases were executed between 2005 and 2015 while excluding sections in which leases were executed between 2002 and 2004 (for the most part, our lease data start in 2002). This restriction allows us to eliminate places where leasing was likely motivated by a desire to exploit formations other than the Haynesville (such as the Cotton Valley formation, which sits above the Haynesville). Third, we drop sections that had wells with production in 2006. Fourth, we drop sections where a non-Haynesville well was drilled after 2000. These last two restrictions allow us to eliminate cases where non-Haynesville activity may have led land to already be held by production prior to Haynesville drilling.

Table 1: Section-level summary statistics

Variable	Obs	Mean	Std. Dev.	P5	P50	P95
Section acres	1195	641.3	13.4	619.2	642.3	661.3
Year first lease starts	1195	2006.5	1.4	2005	2006	2009
Year first lease expires	1195	2009.5	1.5	2008	2009	2012
Number of Hay. wells	733	2.2	2.6	1	1	7
Year of first Hay. spud	733	2009.7	1.1	2008	2010	2011

Figure 2: Map of Haynesville units

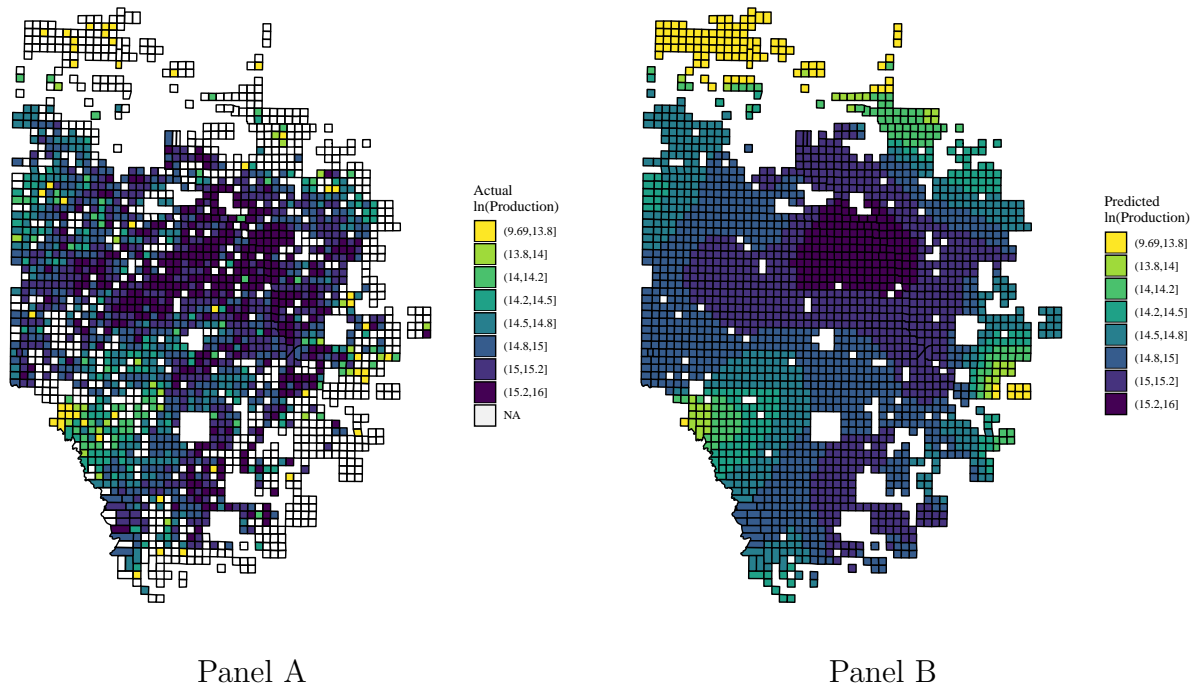


Note: Panel A is a map of Haynesville units, where units that are in the sample are colored dark. The red rectangle is the outline of the map in figure 8. Panel B is a map of Haynesville units, with units colored by how many Haynesville wells were drilled.

of leases have extension clauses, with most extensions lasting 24 months. Extensions require the payment of an additional bonus, but these bonuses are not usually observed in the lease documents. Leases range from less than 0.2 acres to more than 160 acres, with a mean of about 42 acres.⁷ We find that most lessees have small shares, and the HHI for lessee acreage

⁷We have also calculated lease summary statistics weighted by acreage, finding that the distribution of lease characteristics does not change substantially relative to what is presented in table 2.

Figure 3: Haynesville unit productivity



Note: Panel A is a map of Haynesville units, where units are colored by average productivity of wells in the section. Panel B is a map of predicted production after geographic smoothing, discussed in section 6.

is 0.056. The concentration of operators over these leases is somewhat higher at 0.172.⁸

Table 2: Summary statistics for leases

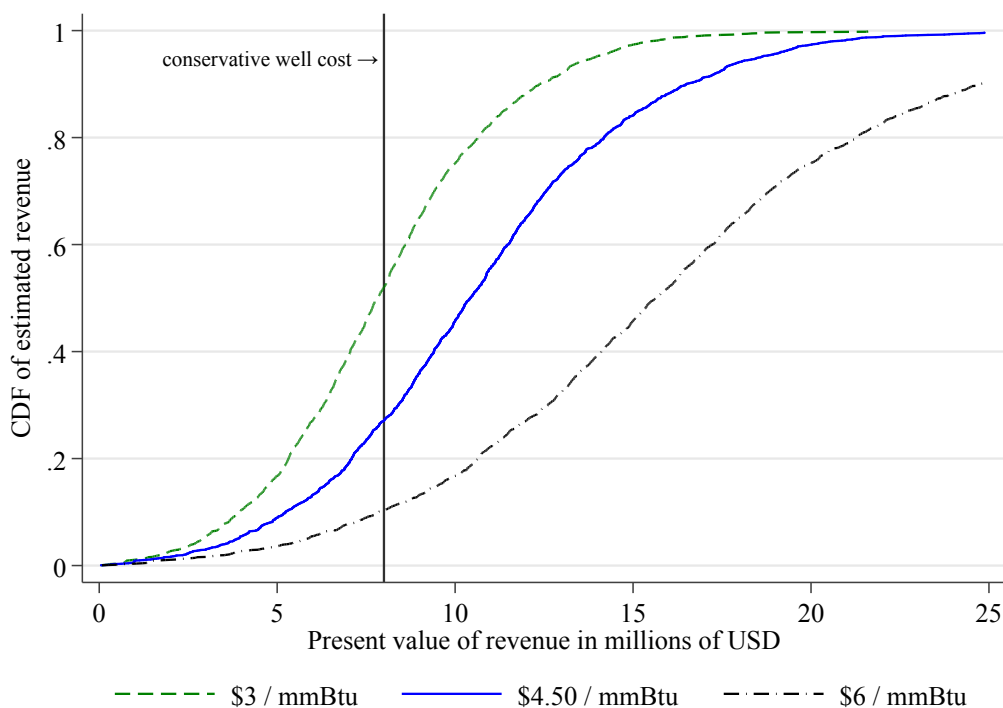
Variable	Obs	Mean	Std. Dev.	P5	P50	P95
Year lease starts	34904	2008.4	1.7	2005	2008	2011
Year lease ends	34904	2011.5	1.8	2008	2012	2014
Primary term length (months)	34904	37.2	6.3	36	36	60
Royalty rate	27339	23	3.4	18.8	25	25
Percent with extension clause	34824	77.7	41.6	0	100	100
Extension length (months)	27055	24.1	2.9	24	24	24
Area in acres	34725	42	153.4	.2	5	162.5

Using data from all Haynesville wells, we examine the distribution of productivity. To construct measures of productivity, we use monthly well-level production data and non-linear least squares to estimate well-level production decline parameters. Our parameterization is

⁸We discuss lessee and operator shares further in appendix C.

based on Patzek et al. (2013) who derive a production decline functional form for shale gas formation wells.⁹ We use the estimated parameters to predict well-level production over time (and beyond our observed data) and then to calculate the present value of each well's total lifetime cumulative production.¹⁰ Figure 3, panel (A) displays a map of these values, averaged at the section-level. We estimate that the median Haynesville well will produce a total discounted quantity of 3.5 trillion British thermal units (3.5 million mmBtu), and that the 25th and 75th percentile wells will produce 2.6 and 4.4 million mmBtu, respectively.¹¹

Figure 4: CDF of present value of lifetime revenues for Haynesville wells



These production calculations suggest that a significant fraction of wells drilled in the Haynesville were unprofitable. In figure 4, we plot the cdf of expected revenues given natural gas prices of \$3, \$4.50, and \$6/mmBtu. The vertical line shows the intersection of our revenue

⁹Appendix G discusses the details of this procedure.

¹⁰We use an annual discount factor of approximately 0.9, consistent with Kellogg (2014).

¹¹One million Btu (mmBtu) is the gas industry's standard unit for natural gas prices. One mmBtu is the energy content of roughly 1,037 thousand cubic feet of natural gas at standard temperature and pressure.

cdfs with a conservative well cost of \$8 million.¹² During the most intense period of drilling in the Haynesville, natural gas prices hovered between \$2 and \$5/mmBtu. At \$4.50/mmBtu, more than 25% of wells drilled appear to be ex-post unprofitable. At \$3/mmBtu, that figure rises to over 50%.¹³

3.2 Empirical evidence on primary terms

To study the role of primary term expirations in motivating what appears to be widespread unprofitable drilling in the Haynesville, we compare the date that the first Haynesville well is spudded in each section to the first date that a lease within the section reaches the end of its primary term (expires). Figure 5 presents a kernel-smoothed distribution of spud timing relative to that expiration date, along with a 95% confidence interval. Figure 6 presents a histogram of the same data. The spike in the density prior to the expiration of the first lease suggests that the expiration is sometimes a binding constraint. Drilling rates are higher in the three month period right before the first lease expires than they are both in the previous three-month period as well as in the following three-month period.¹⁴

Some leases in the Haynesville have a two-year extension clause built in where the firm must pay an additional bonus payment to exercise this option. Figures 5 and 6 show a corresponding spike in drilling roughly two years after the primary term expires. Figure 7 compares sections where the first lease to expire had an extension clause with those that did not. We find that sections with extensions had a less pronounced drilling spike prior to the expiration of the primary term at 36 months and a larger drilling spike prior to the expiration of the extension term at 60 months.

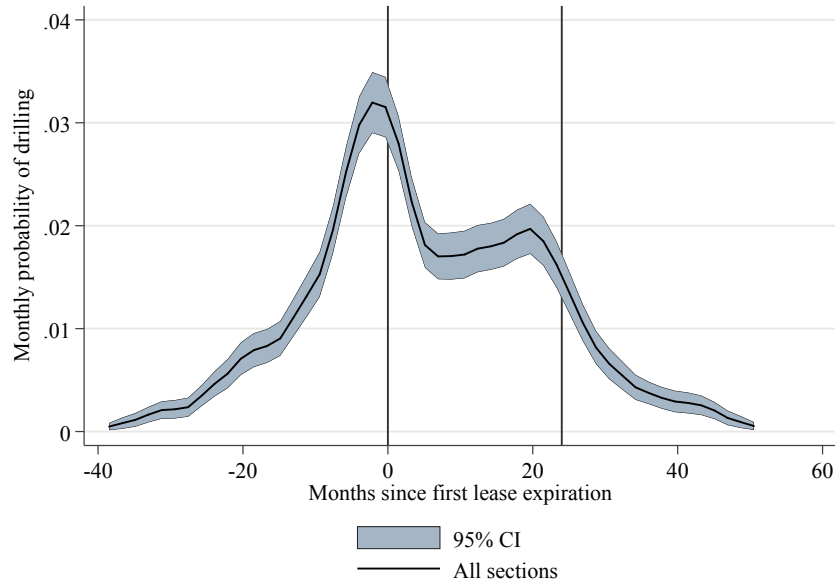
If the primary term is pushing firms to drill a well to hold the lease when they otherwise

¹²Kaiser and Yu (2011) estimate Haynesville drilling and completion costs are \$7 to \$10 million per well. Figures from Berman (2015) suggest that \$8 million per well is a conservative estimate. The median Haynesville well cost reported to the DNR is roughly \$9.4 million.

¹³This criterion is fairly conservative. Because future prices are uncertain, the optimal trigger price for drilling would be higher than that which just covers costs (Kellogg 2014).


¹⁴In appendix C, we use a bunching test to show that the spike in drilling just prior to the primary term expiring is both large and statistically significant.

Figure 5: Kernel-smoothed estimates of the probability of drilling the first Haynesville well in a unit on a given date, relative to the expiration date of the first lease within the unit to expire




Note: Vertical lines are drawn at the date of first lease expiration and two years after first lease expiration.

Figure 6: Histogram of timing of the first Haynesville well drilled in a unit relative to the expiration date of the first lease within the unit to expire




Note: Vertical lines are drawn at the date of first lease expiration and two years after first lease expiration. Bars are 3 months wide.

Figure 7: Smoothed probability estimates comparing sections where the first expiring lease had an approximately two-year extension clause versus sections where the first expiring lease did not have any extension clause



Note: The estimates show the probability of drilling the first Haynesville well in a unit on a given date, relative to the expiration date of the first lease within the unit to expire. Vertical lines are drawn at the date of first lease expiration and two years after first lease expiration.

Figure 8: An example of drilling patterns in the Haynesville Shale.



Note: Map produced using data from Louisiana’s Department of Natural Resources SONRIS. White dots are wellheads, and black lines are the approximate horizontal well path. Units are colored by unit operator. The red rectangle in Panel (A) of figure 2 plots where in the Haynesville unit map this example is located.

would not drill, we would expect that many units would have only a single well for an extended period of time. Figure 8 shows a March, 2017 snapshot of well laterals and pooling units for a selected portion of the Haynesville. Most of these wells were drilled prior to 2013. The fact that so many units still had only one well drilled nearly five years after the drilling boom subsided suggests that drilling was primarily motivated by the possibility of drilling additional future wells, rather than the prospect of immediate profits.

The variation in shading in figure 8 represents different unit operators. We find that operators often to have control of multiple continuous units, which suggests that the drilling patterns are not being driven by externalities like common pool inefficiencies or information spillovers. We examine this possibility further in appendix C, where we show that there is a spike in drilling prior to primary term expiration regardless of whether or not the unit operator controls nearby units.

Finally, we examine the hypothesis that primary terms are less likely to be binding if the firm expects drilling to be highly profitable. In figure 9, we compare sections in the highest tercile of expected discounted cumulative production with those in the lowest tercile of production.¹⁵ Similarly, in figure 10, we use number of wells as a proxy for productivity and compare sections where multiple wells were drilled with ones where only a single well was drilled. In both figures, we find that higher productivity sections tend to have relatively more drilling early in the primary term and less drilling just before the primary term expires.

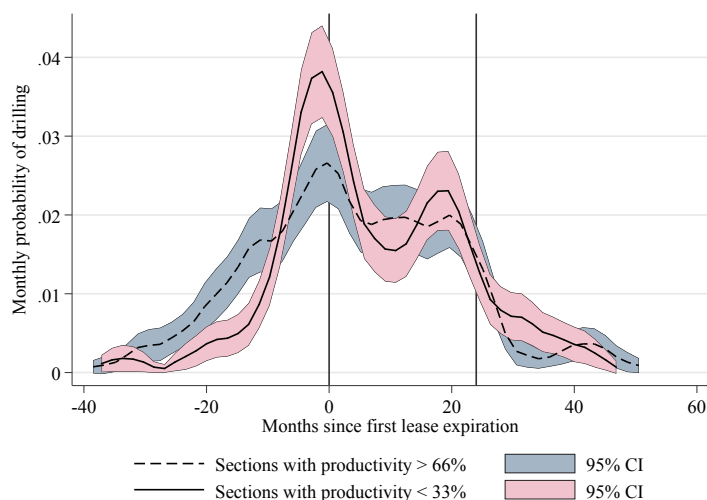
Appendix C contains additional descriptive analysis, including a discussion of how drilling costs and production vary as a function of drilling time, as well as a discussion of how the distribution of primary term expiration dates affects the timing of drilling.

4 An analytic model of oil and gas lease design

In the previous section, we found that primary terms are potentially quite distortionary. In this section, we rationalize these contract features in the context of a principal-agent model


¹⁵We discuss how we estimate section-level production predictions in section 6.

Figure 9: Smoothed probability estimates comparing sections where the first well drilled had expected production in the bottom tercile versus the top tercile



Note: Our calculation of expected production is discussed in section 6. The estimates show the probability of drilling the first Haynesville well in a unit on a given date, relative to the expiration date of the first lease within the unit to expire. Vertical lines are drawn at the date of first lease expiration and two years after first lease expiration.

Figure 10: Smoothed probability estimates comparing sections where there was one well drilled versus sections that had multiple wells drilled



Note: The estimates show the probability of drilling the first Haynesville well in a unit on a given date, relative to the expiration date of the first lease within the unit to expire. Vertical lines are drawn at the date of first lease expiration and two years after first lease expiration.

in which a principal (i.e., the mineral owner) seeks to maximize revenue by contracting with an agent (i.e., the firm) who possesses private information and can take hidden actions. As we will show, these asymmetries drive a wedge between profit maximization and social welfare maximization and lead to contracts similar to those seen in practice.

We model the mineral rights owner as a principal who seeks to contract with a firm, the agent, to extract the natural gas under her land.¹⁶ The central tension in the model lies between the principal’s desire to minimize the firm’s information rent and the need to not overly distort the firm’s incentives to efficiently extract the gas. The information rent rises from the firm’s superior knowledge of the quantity of underground gas reserves and its ability to efficiently extract them.

How should the principal write the contract, assuming that she has the power to set the contractual terms?¹⁷ We assume that the contract can be contingent on two ex-post outcomes that are driven by the firm’s choices: (1) the date at which the well is completed; and (2) the quantity of gas extracted. In practice, both outcomes are easily observable and verifiable, since drilling a well creates an obvious surface disturbance (and requires government permits) and because output is metered and reported to regulatory authorities. We treat other dimensions of the firm’s extraction “effort”, such as the quantity and quality

¹⁶We assume for now that the mineral owner is proposing contracts to a single firm rather than to a set of competing firms, in the spirit of Laffont and Tirole (1986). Relaxing this assumption (which we will do in future iterations of this paper) should not change our qualitative conclusions regarding the structure of the optimal contract. Laffont and Tirole (1987) and McAfee and McMillan (1987) extend the Laffont and Tirole (1986) model to include competition, showing that the contract structure is unchanged (though the contract selected by the winning firm becomes “flatter” in expectation as competition increases). Similar results hold with increased competition in Board (2007).

¹⁷As DeMarzo et al. (2005) explain, this assumption is not innocuous. A contingent contract is optimal for the principal if she holds all the bargaining power. If the agent, however, can make a take it or leave it offer, the equilibrium contract will involve only a flat cash payment to the principal (and overall lower expected revenue for the principal). Our model is therefore more closely aligned with state and federal oil and gas auctions—where the governments involved clearly hold the power to propose terms—than with onshore contracts between private mineral owners and extraction firms. That said, in the private minerals setting firms do not seemingly have the power to make take it or leave it offers either. Anecdotally, even if the firm makes the initial offer, the final contract is the outcome of a negotiation process in which the mineral owner can propose counter-offers. Moreover, the mineral owner’s outside option (in Louisiana) is to be a working interest owner in the unit, which entitles her to a contingent payment that takes the form of a call option. Thus, even in private mineral settings where contract terms are negotiated, we should not be surprised to observe contract provisions that are qualitatively consistent with mineral owners holding bargaining power.

of its fracking inputs or its engineering design expenditure, as non-contractible. Given these contractibility assumptions, our model then combines features from Laffont and Tirole’s (1986) static model of using ex-post cost observability to write procurement contracts and Board’s (2007) model in which the principal sells a development option to an agent, where the execution date is observable but ex-post profits and revenue are not. As a result, the optimal contract we derive involves both revenue sharing (akin to the cost sharing in Laffont and Tirole 1986) and a payment made at the time of well completion (similar to the payment from the agent to the principal upon the execution of the option in Board 2007, except our payment is from the principal to the agent).

4.1 Setup

The key objects in our model are as follows:

- θ denotes the productivity of the underground natural gas resource should the firm drill and complete a well. θ is known by the firm but not by the principal.¹⁸
- $F(\theta)$ denotes the principal’s rational beliefs about the distribution of θ . $F(\theta)$ has support on $[0, \bar{\theta}]$.
- Time is discrete and denoted by $t \in \{0, \dots, T\}$, where T is possibly infinite. The contract between the principal and the firm is set at $t = 0$, and then starting at $t = 1$ the firm can decide whether to execute the option to drill and complete a well. The principal observes the period in which drilling and completion occur. Only one well may be drilled on the lease.
- $e \in \mathbb{R}^+$ denotes completion “effort”. For instance, e may represent the volume of fracking fluid that is used to frac the well or engineering expenditures on well design. e is not observable to the principal.

¹⁸Allowing for incomplete information on the part of the firm will not affect the paper’s results, so long as the firm is better-informed than the principal.

- The cost of drilling and completing the well is given by $c_0 + c_1e$, where c_0 and c_1 are strictly positive scalars that are common knowledge.
- $y = g(e)\theta(1 + \varepsilon)$ denotes the volume of natural gas extracted if the firm drills and completes a well with effort e . y is observed by both the principal and firm, and for simplicity we assume that y is completely realized in the same period that the well is drilled and completed.¹⁹ The production function $g(e)$ is common knowledge and has the properties that $g : \mathbb{R}^+ \rightarrow [0, 1)$, $g(0) = 0$, $g' > 0$, $\lim_{e \rightarrow 0^+} g'(e) \rightarrow \infty$, and $g'' < 0$. Thus, the function g maps completion effort into a recovery ratio, where complete recovery is never possible even with infinite effort. The variable ε is a mean-zero disturbance that is unknown prior to drilling by both the principal and firm. ε is orthogonal to e and θ , it has support on $[-1, \infty)$, and its distribution function $\Lambda(\varepsilon)$ is common knowledge.
- The oil price at time t is denoted P_t and is common knowledge. The oil price evolves stochastically via a process that is common knowledge and has the property that P_t is bounded above. P^t denotes the entire history of prices from time 0 through t .
- The common per-period discount factor is given by δ . Both the principal and firm are risk neutral.

Both the principal and the firm seek to maximize the expected present value of their respective cash flows. At $t = 0$, the principal can offer a menu of contracts to the firm; the firm must then choose one such contract or decline entirely (yielding a payoff of 0). The contracts can specify transfers that are contingent on observed extraction y , the time at which drilling takes place, and the history of prices up to the time of drilling. After the firm

¹⁹In reality of course, gas is produced over many months and years subsequent to drilling. y can be thought of as the present discounted volume of production over the life of the well, and accordingly P_t can be thought of as the weighted average natural gas price that is expected to prevail over the productive life of a well that is drilled and completed at t . This formulation assumes that well production is not affected by gas price realizations subsequent to drilling and completion, consistent with theory and evidence in Anderson et al. (2018).

selects its contract, it then faces an optimal stopping problem regarding when to drill the well, and when it drills it must choose its effort e .

Section 4.2 below discusses the optimal contract implied by our model. We relegate the derivation of this contract—which draws heavily from Laffont and Tirole (1986) and Board (2007)—to appendix D.

4.2 The revenue-optimal mineral lease

Appendix D shows that the mineral owner’s expected revenue-optimal lease can be implemented by offering the firm a menu of contracts that include a contingent payment that is affine in gas revenue. Specifically, the optimal contingent payment z that is paid from the firm to the principal at well completion time τ , given by equation 21 in appendix D, takes the form:

$$z(\theta, P_\tau y) = z_0(\theta)P_\tau + z_1(\theta)P_\tau y, \tag{1}$$

where $z_0(\theta)$ is a negatively-valued, increasing function of θ , and $z_1(\theta)$ is a positively-valued, decreasing function of θ . Both $z_0(\theta)$ and $z_1(\theta)$ are zero for $\bar{\theta}$, reflecting the standard intuition that the incentives for the highest type are not distorted.

The intuition for the tax on production revenue—i.e., a royalty—in the right-most term of equation (1) flows from the linkage principle (Milgrom and Weber 1982; Riley 1988; DeMarzo et al. 2005). Because production revenue is correlated with the firm’s private information θ , the royalty mitigates the firm’s ability to earn information rent by compressing the distribution of payoffs across types. A 100% royalty is not desirable, however, because such a confiscatory tax would result in the drilling option never being exercised. The optimal royalty therefore strikes a balance between reducing the firm’s information rent and minimizing distortions to the firm’s effort.

In standard mechanism design problems of this type, such as Laffont and Tirole (1986), “effort” is one-dimensional and unobservable. In oil and gas drilling, however, it is useful

to (loosely) think of “effort” as having two dimensions: the decision of when to drill and the decision of how much labor, capital, and material to invest in the well. The former is straightforward for the mineral owner to observe and contract on, while the latter is not. The essence of our result in equation (1) is that, given that the mineral owner wants to tax natural gas production, she can mitigate the resulting distortions to observable dimensions of the firm’s effort by subsidizing them. Specifically, she can subsidize the drilling of the well by paying $z_0(\theta)$ to the firm. In the extreme, if all dimensions of effort were observable, the optimal mechanism would call for the principal to pay the full cost of the well and then receive 100% of the production revenue.

It remains to reconcile the optimal implementation derived above with contracts observed in practice, which feature a primary term rather than a drilling subsidy. Per intuition from Board (2007), the royalty and subsidy in the optimal mechanism act as a Pigouvian tax and subsidy, in that they align the firm’s incentives with the principal’s. A primary term is qualitatively similar to a drilling subsidy, in the sense that both provide the firm with an incentive to drill sooner than it otherwise would, given the royalty. If the future path of natural gas prices were certain at $t = 0$, the two policies would be able to achieve the same outcome. But gas prices are of course stochastic, and per the intuition given by Weitzman (1974) and Kaplow and Shavell (2002), we know that a “quantity policy” such as a primary term will result in an expected utility loss relative to a “price policy” that is an optimal Pigouvian subsidy. In particular, the primary term will result in drilling occurring too soon if future prices are lower than expected, and it may result in drilling occurring too late if future prices are higher than expected.

So why use a primary term rather than the subsidy? In our view, the most likely candidate for the discrepancy between our theory and practice is that mineral owners may be liquidity constrained in their ability to transfer cash to the firm. Indeed, one reason why mineral owners contract with oil and gas firms in the first place is to address their inability to finance resource extraction themselves. Liquidity constraints clearly do not apply to federal

offshore leases, but in that case there may be political and budgetary constraints to having the government fund oil and gas drilling, particularly when the decision-making is under the control of the firm.

It remains to examine whether using a primary term rather than a drilling subsidy can improve expected revenue for the principal. Because modeling the expected revenue from a lease with a primary term is not analytically tractable, we turn to a computational model of the firm’s drilling decisions, its profits, and the mineral owner’s expected revenue under alternative lease designs. This computational approach will also allow us to incorporate several other features of the Haynesville into our analysis, including the ability of firms to hold an entire unit—and therefore the option to drill multiple future wells—by completing a single well.

5 Computational model

This section summarizes our computational model of the firm’s drilling timing decision and the resulting value that accrues to both the firm and the mineral owner, depending on the lease terms. Additional detail is provided in appendix E.

As with the analytic model, the extraction firm has a type θ drawn from a distribution $F(\theta)$. While the firm knows θ , the mineral owner only knows $F(\theta)$. At the beginning of the game ($t = 0$), the owner makes a take-it-or-leave-it lease offer to the firm. The contract includes a primary term $\bar{T} \in \{\mathbb{N}, \infty\}$ that denotes in which period the lease expires, a royalty rate $k_1 \in [0, 1]$ that denotes what fraction of the revenues will be paid to the owner, and a bonus $R_1 \in \mathbb{R}$ that is paid to the owner when the firm signs the lease. The contract may also include a subsidy for drilling $S_1 \in \mathbb{R}$. We notate the vector of lease characteristics $\chi_1 = \{\bar{T}, k_1, R_1, S_1\}$. If the firm accepts the lease offer, it pays the bonus R_1 in period $t = 0$ and the game continues.

In each period $t = 1$ to $t = \bar{T}$, the firm chooses to drill one well or not. For now, we

assume that the lease can accommodate only a single well, although we generalize the model to multiple wells later. The payoff to drilling is affected by the gas price P_t and drilling cost C_t which evolve according to a common knowledge first-order Markov process. If the firm drills, it pays drilling costs, collects the subsidy S_1 (if applicable), extracts the net present total value of the well, and pays royalties to the owner. If the firm does not drill, the game continues to period $t + 1$.

When the lease is about to expire at $t = \bar{T}$, the owner makes a take-it-or-leave-it offer of an extension with contract terms $\chi_2 = \{k_2, R_2, S_2\}$. For computational simplicity, we assume that the extension term is infinite. We also assume in our simulations that $k_2 = k_1$ and $S_2 = S_1$, reflecting common practice in the Haynesville. Faced with this new contract, the firm chooses either to drill, abandon the lease, or pay the extension bonus R_2 to extend the lease. If the firm pays to extend, then for periods $t \geq \bar{T} + 1$ the firm solves an infinite-horizon optimal stopping problem in which it must decide when to drill.

Table 3 summarizes the timing of the model and the action space of the firm, and table 4 summarizes the per-period payoffs to the firm and the owner for each action. For each action a and for all periods $t \geq 1$, we include an action-specific shock $\nu_{t,\theta}^a$ to the per-period payoff of the firm. We assume that the firm's shocks $\nu_{t,\theta}^a$ are drawn from an i.i.d. type 1 extreme value distribution with scale parameter σ_ν . These idiosyncratic shocks capture unexpected transitory drivers of drilling behavior (such as drilling and completion contractor availability) and have the effect of smoothing the model's predicted time path of drilling.²⁰

The game is solved computationally via backward induction. During the extension period $t > \bar{T}$, the firm makes optimal drilling timing decisions under an infinite horizon given contract terms χ_2 . In period $t = \bar{T}$, the owner forecasts the firm's future drilling decisions and chooses the contract terms χ_2 , including the extension bonus R_2 , that maximizes the

²⁰For all periods $t \notin \{0, \bar{T}\}$, the possible actions are to drill ($a = 1$) or not drill ($a = 0$). For $t = \bar{T}$, we use a nested setup: First the agent chooses whether to drill ($a = 1$, with additive shock $\nu_{t,\theta}^1$) or not. Then conditional on not drilling, the firm chooses whether to extend or abandon. We assume that the firm receives the same additive shock $\nu_{t,\theta}^0$ for both abandoning the lease and for continuing the lease. More details are in appendix E.

Table 3: Timing and order of actions in the computational model

Time period	Lease stage	Order and actions
$t = 0$	$i = 0$	Nature draws prices and costs Owner chooses initial lease terms χ_1 , including \bar{T} Firm chooses to accept contract χ_1 or reject
$1 \leq t \leq \bar{T} - 1$	$i = 1$	Nature draws prices and costs Firm chooses to drill or wait
$t = \bar{T}$	$i = 1$	Nature draws prices and costs Owner chooses extension lease terms χ_2 Firm chooses to drill, pay bonus to continue, or abandon
$t \geq \bar{T} + 1$	$i = 2$	Nature draws prices and costs Firm chooses to drill or wait

Table 4: Table of the firm's action choices and corresponding payoffs

Time periods	Firm action	Firm payoff	Owner payoff
$t = 0$	Sign lease	$-R_1$	R_1
	Do not sign	0	0
$t \geq 1, t \neq \bar{T}$	Drill	$(1 - k_i)P_t\theta - C_t + S_1 + \nu_{t,\theta}^1$	$k_iP_t\theta - S_1$
	Continue	$\nu_{t,\theta}^0$	0
$t = \bar{T}$	Drill	$(1 - k_1)P_t\theta - C_t + S_1 + \nu_{t,\theta}^1$	$k_1P_t\theta - S_1$
	Extend	$-R_2 + \nu_{t,\theta}^0$	R_2
	Abandon	$\nu_{t,\theta}^0$	0

owner’s payoff.²¹ During the primary term periods $t < \bar{T}$, the firm forecasts the extension contract terms χ_2 (including the extension bonus R_2) that the owner will offer in $t = \bar{T}$ and incorporates that into the costs of waiting. In the initial time period $t = 0$ the owner anticipates future drilling decisions and sets contract terms χ_1 , including the bonus R_1 , that maximize its profits. Because we assume no random shocks in the initial period, there will be a threshold type θ^* such that all types with $\theta \geq \theta^*$ accept the initial contract.

In appendix F, we discuss how we extend the model to the case of multiple wells per unit. In this extension, we allow the firm to drill an additional $N - 1$ wells in either the same period the first well is drilled or any period thereafter. Because drilling the initial well holds the lease by production, so that the firm is no longer subject to lease expiration, the firm may choose to substantially delay the drilling of the $N - 1$ wells relative to the first well.

6 Calibration

In this section we discuss the parameters used in the computational model, including drilling costs, stochastic price processes, and the distribution of productivity $F(\theta)$. Our goal is to calibrate to a representative pooling unit in the Haynesville shale, so we use data from the Haynesville to construct appropriate parameters reflecting market conditions at that time. We use a lease start date of fourth quarter of 2006, as that was the median date of first leasing within our sample of units. A summary of calibrated parameters is near the end of the section in table 5. Additional detail on our data sources is provided in appendix B.

Natural gas prices: We calibrate the stochastic process for natural gas prices P_t using futures prices for delivery at Henry Hub, Louisiana. Rather than use spot or front-month

²¹We simplify this step of the model by assuming that the mineral owner sets χ_2 as though it still faces the original distribution of firm types. In principle, the owner should realize that a high-type firm would likely have already drilled by the time \bar{T} is reached, and a low-type firm would not have accepted the original contract χ_1 . However, allowing the mineral owner to update its beliefs in this manner would result in an intractably complicated state space and a much larger computational burden, since the type distribution at \bar{T} is a function of the complete realized price and drilling cost path during the primary term.

prices, we use prices for delivery at a 12-month horizon because wells produce gas gradually rather than instantaneously, and because the 12-month horizon is the longest at which futures are consistently liquidly traded. We aggregate the raw daily price data to the quarterly level (to match the timing of the model) and deflate it to December, 2014 dollars.²²

Following Kellogg (2014), we assume that P_t follows a first-order Markov process of the following form:

$$\ln P_{t+1} = \ln P_t + \kappa_0^P + \kappa_1^P P_t + \sigma^P \eta_{t+1}^P, \quad (2)$$

where the drift parameters κ_0^P and κ_1^P allow for mean reversion. We assume that price volatility σ^P is constant. We estimate κ_0^P and κ_1^P by regressing $\ln P_{t+1} - \ln P_t$ on P_t , using data from 1993 (when futures prices are first reliably liquid) through 2009 (just before the peak of Haynesville drilling). The parameter estimates are shown in table 5 and imply that the long-run mean natural gas price is \$4.99/mmBtu.

Severance and income taxes: Louisiana, like many states, imposes a severance tax on the production from oil and gas wells. As discussed by Kaiser (2012), the severance tax on Haynesville shale wells is paid after either the well has been producing for two years or the well’s drilling costs have been paid, whichever comes first. Because modeling this process is complex, we follow Gülen et al. (2015) by assuming a tax of 4% on production revenue, and we allow the firm to deduct drilling costs (subject to revenue exceeding costs). The severance tax applies to both the firm’s working interest production and the mineral owner’s royalty share.

Per Gülen et al. (2015), the combined state and federal marginal corporate tax rate in Louisiana is 40.2%. Following Gülen et al. (2015), we treat 50% of drilling and completion expenditures as immediately expensable, while the remainder must be capitalized and depreciated over time. Per Metcalf (2010), we depreciate the capitalized drilling and completion

²²We use the Bureau of Labor Statistics’ Consumer Price Index for all goods less energy, all urban consumers, and not seasonally adjusted. The CPI series ID is CUUR0000SA0LE.

costs over seven years using the double declining balance method. The effective marginal tax rate on drilling and completion costs is then 36.8%, so that income taxes have the effect of modestly discouraging drilling.

We assume that the marginal tax rate of 40.2% also applies to the mineral owner’s royalty income. By equating the mineral owner’s and the firm’s tax rate on revenues, we avoid building into the model any tax advantages for shifting income from the firm to the mineral owner (or vice-versa). Moreover, all bonus payments and drilling subsidies in our simulations can then be interpreted as post-tax payments.²³

Drilling and completion costs: We assume that the well drilling and completion cost C_t is a function of the rig dayrate D_t , which represents the per-day rental rate for a drilling rig and crew. We obtain quarterly data on dayrates from RigData, a private industry intelligence firm, and we deflate these data to December, 2014 dollars. As with the stochastic process for gas prices, we assume that dayrates follow the first-order Markov process given by equation (3):

$$\ln D_{t+1} = \ln D_t + \kappa_0^D + \kappa_1^D D_t + \sigma^D \eta_{t+1}^D. \quad (3)$$

We assume that $\kappa_0^D = \kappa_0^P$ and $\kappa_1^D = \kappa_1^P \bar{D}_t / \bar{P}_t$, so that dayrate mean reversion is proportional to that of natural gas prices.²⁴ These parameters imply that the long-run mean dayrate is \$9,597/day. We assume that the shocks η_{t+1}^D and η_{t+1}^P are each drawn from an i.i.d. bivariate standard normal distribution, with a covariance matrix that we estimate using the residuals of equations (2) and (3).

We then assume that drilling and completion costs C_t are an affine function of D_t , per equation (4):

²³In reality, the marginal tax rate faced by mineral owners is likely to be quite heterogeneous given underlying heterogeneity in incomes from other sources. We also do not model the fact that royalties are tax-advantaged relative to bonus payments, since the firm must capitalize bonuses but can expense royalties.

²⁴ \bar{P}_t and \bar{D}_t denote the average price and dayrate, respectively, over 1993–2009.

$$C_{it} = \alpha_0 + \alpha_1 D_t. \quad (4)$$

We estimate equation (4) using well-level drilling and completion cost reports that unit operators file with the Louisiana DNR for the purpose of determining severance taxes. Accordingly, these data may not be perfect measures of true economic drilling costs; however, we view them as a useful proxy for the purpose of model calibration. We estimate that $\alpha_0 = \$7.1$ million and $\alpha_1 = 159$ days.²⁵ To provide a sense of magnitudes, the average rig dayrate in 2010 was \$15,700/day, so that our projected average drilling and completion cost for 2010 is \$9.6 million.

Operating and gathering costs: We assume that production from each well incurs operating and gathering costs (combined) of \$0.60/mmBtu (gathering costs reflect costs of transporting gas from the Haynesville to the Henry Hub location in southern Louisiana). We obtain this value from Gülen et al. (2015), which calculates costs of \$0.55–\$0.60/mmBtu for most wells in the early years of operation, increasing thereafter as fixed operating costs are spread over declining production. We simplify the approach of Gülen et al. (2015) by assuming that operating and gathering costs are \$0.60/mmBtu throughout the life of the well. We also assume that operating and gathering costs are fully expensable for income tax, but not severance tax, purposes.

Distribution of payoff shocks The model’s unobserved payoff shocks $\nu_{t,\theta}^a$ for each action have a type 1 extreme value distribution with mean zero and scale parameter σ_ν . In a full structural estimation, the magnitude of σ_ν would be pinned down by the responsiveness of drilling to price and dayrate shocks. For now, however, we calibrate σ_ν by interpreting the $\nu_{t,\theta}^a$ as pure drilling cost shocks. We then take σ_ν to be the root mean squared error from

²⁵Because the drilling rig dayrate is one of several service costs that scale with the time spent drilling and completing a well, our estimate of α_1 is inflated relative to actual drilling and completion times. Our projection in equation 4 essentially treats variation in rig dayrates as an index for variation in overall drilling and completion service costs.

the regression of drilling costs on rig dayrates. The resulting value is $\sigma_\nu = \$1.29$ million.²⁶

Distribution of θ : We use our well-level lifetime cumulative production estimates, discussed in section 3, to estimate the distribution of θ , the firm’s expectation of the net present value of total well production. We estimate this distribution, $F(\theta)$, in two steps. First, we use a partially linear model, discussed in detail in appendix H, to generate predictions of log well productivity that are geographically smoothed and account for a flexible time trend in productivity using the difference estimator described in Robinson (1988). The productivity smoothing procedure is optimized for out-of-sample prediction using leave-one-out cross-validation.

We adjust for the change in productivity between the earliest Haynesville wells (the median unit’s first lease signing is 2006 Q4) and the sample average.²⁷ Specifically, we adjust our smoothed productivity estimates by evaluating them at the earliest observed completion date of October 10, 2007, which is 0.97 log points lower than the mean time effect. The output of this procedure is a set of unit-level expected log gas production values evaluated at October 10, 2007 productivity levels, which we denote as $\ln(\tilde{\theta})$ and map in figure 3, panel (B).

In our second step, we use the estimated $\ln(\tilde{\theta})$ to estimate the parameters of the distribution $F(\theta)$. We assume that θ is distributed log normally with parameters μ_θ and σ_θ . We use the standard deviation of $\ln(\tilde{\theta})$, which we find to be 0.53, as our measure of σ_θ . This calculation assumes that the mineral owner knows the distribution of predicted section-level productivity across the Haynesville play but does not know where the productivity of its own parcel falls within this distribution. That is, we are assuming that the owner of each parcel i knows the distribution mapped in figure 3, panel (B) but does not know θ_i . This approach may over-estimate mineral owners’ uncertainty over $\tilde{\theta}$ to the extent that owners

²⁶We reduce the scale parameter by 40.2% so that it is in post-income tax terms.

²⁷Our modeling assumes that mineral owners and firms did not foresee the productivity growth in the Haynesville that occurred between the periods when most leases were signed and when most drilling occurred. Our model will under-estimate total lease values to the extent that productivity growth was anticipated. Moreover, throughout the firm’s simulated drilling problem, we hold productivity constant at its initial value.

Table 5: Summary of parameters for calibration

Parameter	Notation	Value	Source
Drilling Costs			
Intercept	C_t	LA DNR reported well cost data	
Dayrate coefficient	α_0 α_1	\$7.08 million 158.97 days	
State transitions			
Price drift constant	$\{P_t, D_t\}$ κ_0^P	0.0058	Henry Hub spot prices and RigData dayrates
Price drift linear term	κ_1^P	-0.0022	
Price volatility	σ^P	0.102	
Dayrate drift constant	κ_0^D	0.0058	
Dayrate drift linear term	κ_1^D	-1.05×10^{-6}	
Dayrate volatility	σ^D	0.093	
Price - dayrate correlation	ρ	0.373	
Action-specific shocks			
Scale parameter	ν^a σ_ν	\$1.29 million	RMSE from well cost regression
Well productivity			
Mean of $\ln(\theta)$	$\theta \sim F(\theta)$ μ_θ	mmBtu 13.83	Analysis of Haynesville production data
SD of $\ln(\theta)$	σ_θ	0.53	
Taxes			
Severance taxes		4% of revenues	Gülen et al. (2015)
Federal and state income taxes		40.2% of revenues	

have at least some rough knowledge of whether their parcel lies on a relatively productive or unproductive part of the Haynesville. Use of smaller values for σ_θ would imply smaller optimal royalties in our simulations below.

We calibrate μ_θ so that the mineral owner’s expected value of θ is equal to the median of the $\tilde{\theta}$. We therefore calculate μ_θ as the median value of $\ln(\tilde{\theta})$, less $\sigma_\theta^2/2$, yielding a value of 13.83, equivalent to 1.01 million mmBtu. At the 2006 Q4 price of \$10.22/mmBtu, this volume of gas is worth \$10.4 million.

7 Results

The computational model presented in section 5 and calibrated per section 6 provides us with a tool to evaluate how drilling outcomes, the firm’s profits, and the mineral owner’s profits vary with three key lease contract variables: the royalty, primary term, and drilling subsidy. For each set of lease conditions we examine, we assume that the mineral owner offers a take-it-or-leave-it bonus payment that is expected revenue optimal conditional on the calibrated parameters, lease contract terms, and the state variables (the gas price and rig dayrate) at the time of lease signing.

7.1 Optimal contracts for simple units with one well

We begin by studying the mineral owner’s optimal contract for a simple drilling unit on which only a single well can be drilled. We search for the optimal contract by computing the owner’s payoff over a grid of possible royalty rates (in increments of 5%) and primary terms $\{0.25, 0.5, 0.75, \dots, 2.5, 3, \dots, 5, 6, \dots, 15, \infty \text{ years}\}$. When the initial gas price is \$10.22/mmBtu (the price prevailing in the fourth quarter of 2006), we find that the mineral owner achieves its highest expected payoff with a 60% royalty rate combined with a 3.5-year primary term.²⁸

This royalty rate is quite a bit higher than the 20–25% royalty rates typically observed

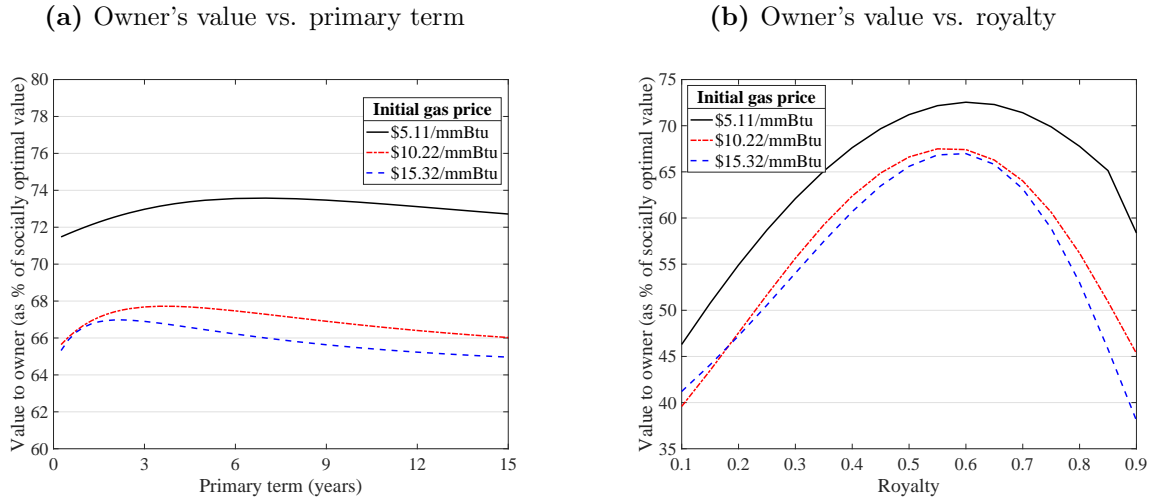
²⁸For reference, we find that the optimal bonus payment is \$199,000.

in the Haynesville Shale. We believe that there are at least four factors driving our optimal royalty rate higher. First, we do not explicitly model unobserved firm effort such as fracking input intensity; in reality, the royalty distorts effort, thus making higher royalty rates less desirable from the perspective of the mineral owner. Second, the optimal royalty rate is closely related to the mineral owner’s uncertainty with respect to productivity. As discussed in the calibration section, our value of σ_θ is drawn from the dispersion of predicted production across the entire Haynesville Shale. If mineral owners have at least a rough idea of whether they are in a relatively productive or unproductive part of the shale (even at a parish level), actual information asymmetry is less than what we are modeling, which will lead to a lower optimal royalty. Third, mineral owners may be risk-averse relative to firms, which will lead them to prefer larger bonuses over larger royalties. Finally, our optimal bonus calculations assume that there is only one firm to whom to offer a lease. In the presence of competition in the leasing market, bonus auctions become more effective at reducing the firms’ information rents, so that the optimal royalty is reduced.

A central puzzle motivating this paper is that leases include primary terms even though these terms induce distortions in firm behavior. The analytical model shows that a cost subsidy allows the mineral owner to improve its payoffs by countering some of the distortion induced by the royalty. However, a primary term is merely a coarse approximation of the cost subsidy. In figure 11, panel (a), we fix the royalty rate at 60%, vary the length of the primary term, and plot the mineral owner’s payoff as a percentage of the socially optimal total value. Each line represents a different assumption about the prevailing gas price at the time the lease is signed. The optimal primary term is the point at which these lines reach a maximum.

One immediate result is that the optimal primary term is finite in all cases shown. That is, the presence of a properly-chosen primary term can indeed make the mineral owner better off. The red dash-dot line shows that the optimal primary term is 3.5 years when the initial price is \$10.22/mmBtu. When the initial price is instead 50% higher (\$15.32/mmBtu), the

Figure 11: Expected value to the mineral owner as a function of lease terms



Note: In panel (a), the royalty is held fixed at 60%. In panel (b), the primary term is held fixed at 3.5 years. The drilling subsidy paid by the mineral owner is \$0. All expected values are shown as a percentage of the expected value from the socially optimal drilling program. See text for details.

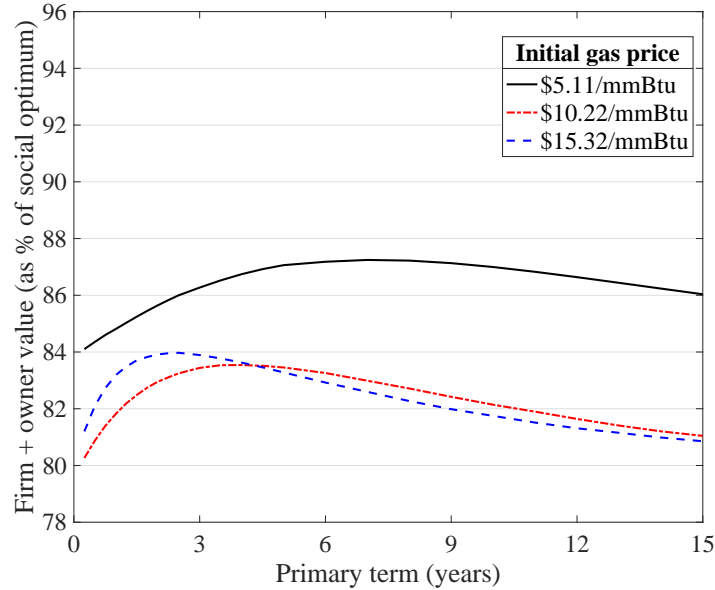
optimal primary term is just 2.25 years, as shown by the dashed blue line. Conversely, the optimal primary term is 7 years when the initial gas price is 50% lower (solid black line). Intuitively, when the initial gas price is low, drilling early in the lease is likely to be socially undesirable, so the optimal primary term is relatively long.

Panel (b) of figure 11 shows how the mineral owner's expected revenue varies with the royalty rate, holding the primary term fixed at 3.5 years. At all three initial gas prices, the owner's revenue is very sensitive to the royalty rate, and the red dash-dot line shows the optimality of the 60% royalty at the baseline initial price of \$10.22/mmBtu. The optimal royalty decreases modestly with the initial gas price, reflecting the fact that the dispersion in the firms' profit types is decreasing with the initial gas price (because all firms have the same costs, the dispersion in profits that is induced by dispersion in expected revenue is greater when the gas price is low).²⁹

While the mineral owner appears to be better off with a primary term, what about the

²⁹A countervailing effect is that at high gas prices, substantial drilling delays are induced only at higher royalties. On net then, the optimal royalty is not particularly sensitive to the natural gas price.

Figure 12: Sum of the mineral owner’s and firm’s expected values as a function of the primary term

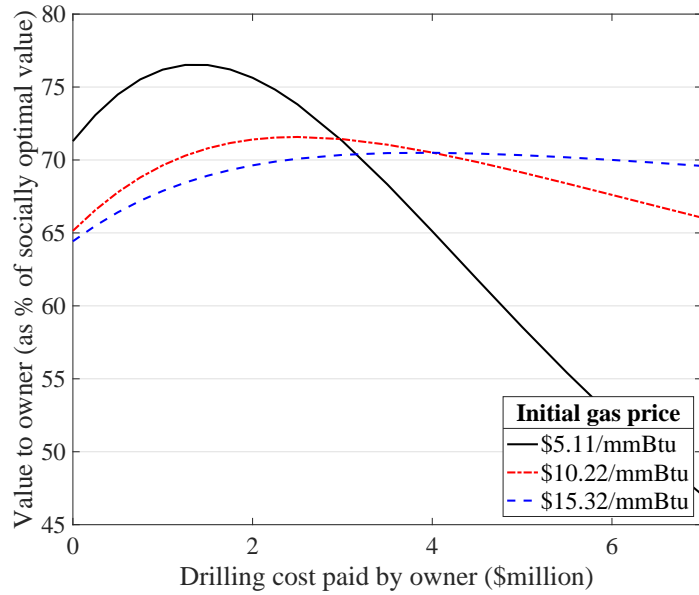


Note: Over all primary terms, the royalty is held fixed at 60%. The drilling subsidy paid by the mineral owner is \$0. All expected values are shown as a percentage of the expected value from the socially optimal drilling program. See text for details.

overall surplus? In figure 12, we plot the sum of the mineral owner’s and firm’s payoffs, which is the total expected surplus from the lease. The surplus-maximizing primary term is finite and quite similar to the mineral owner’s value-maximizing primary term. This figure also reveals that at the baseline initial price of \$10.22/mmBtu, the overall welfare loss generated by the owner’s optimal lease contract is approximately 16% of social value.

The primary term improves the owner’s payoff in the presence of a royalty payment, but how does it compare to the optimal drilling subsidy? Figure 13 shows how the mineral owner’s payoff varies with the drilling cost subsidy, for a lease with a royalty of 60% and an infinite primary term. For each initial gas price, the optimal drilling subsidy outperforms the optimal primary term from figure 11, panel (a). To illustrate the magnitudes at stake, consider the baseline \$10.22/mmBtu case. The mineral owner earns (in expectation) 71.6% of the socially optimal surplus using the optimal \$2.5 million drilling cost subsidy, relative

Figure 13: Expected value to the mineral owner as a function of the drilling subsidy



Note: Over all drilling subsidies, the royalty is held fixed at 60%. The primary term is infinite. All expected values are shown as a percentage of the expected value from the socially optimal drilling program. See text for details.

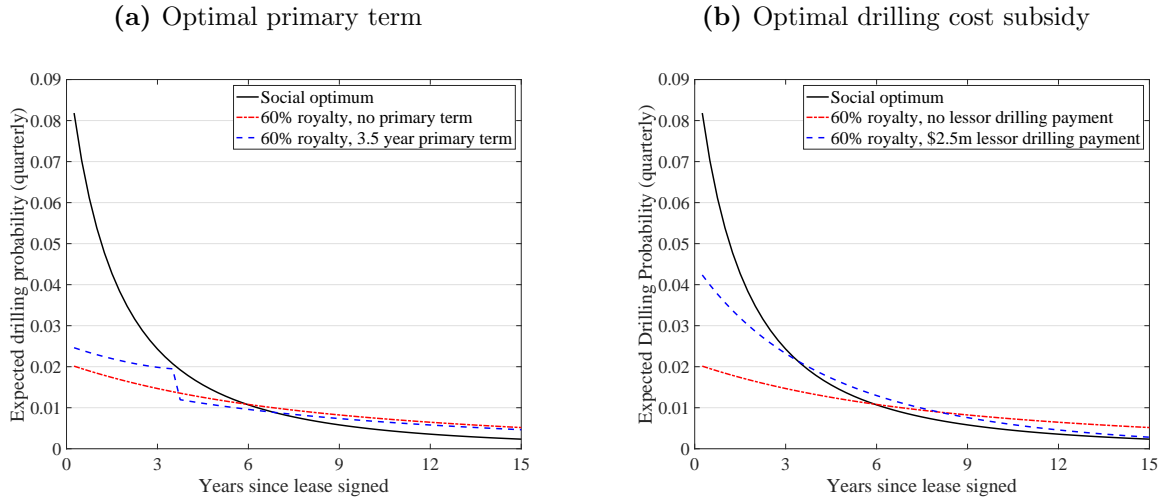
to 65.1% when there is only a royalty.³⁰ This gain corresponds to \$225,000. Under the optimal primary term, the mineral owner earns 67.7% of the socially optimal surplus; the gain relative to the royalty-only contract is \$90,000. The optimal primary term therefore achieves only 40% of the gains realized by the optimal subsidy.

In this simplified model, the timing of the drilling decision is the primary distortion driven by the lease contract.³¹ We can compare drilling timing under various contract terms to the socially optimal drilling program. In figure 14, panel (a), we plot the probability that the unit is drilled in a given quarter. We fix the initial natural gas price to \$10.22/mmBtu and simulate three cases: the social optimum, a 60% royalty-only contract, and a 60% royalty contract with a the optimal 3.5-year primary term. Comparing the social optimum to the royalty-only contract highlights the distortion that leads to sub-optimally late drilling. While

³⁰The optimal bonus associated with the \$2.5 million drilling cost subsidy contract is \$1.67 million.

³¹The other distortion is in the productivity cutoff below which the firm will not accept the offered contract.

Figure 14: Drilling probabilities as a function of lease terms



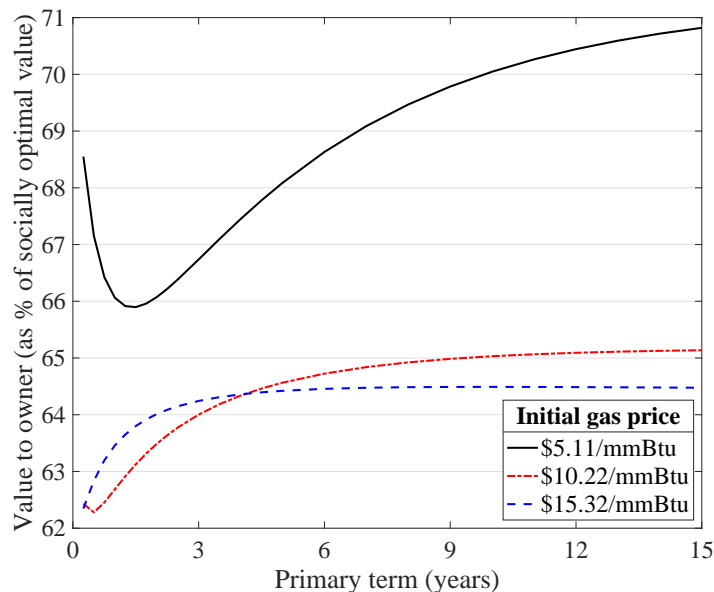
the primary term pulls drilling forward relative to the royalty-only case, it does not very well approximate the drilling probability of the social optimum. Instead, drilling activity moves forward to the primary term but drops off sharply in the periods after expiration. In contrast, panel (b) of figure 14 replaces the primary term with the optimal drilling subsidy of \$2.5 million. The drilling probabilities are still delayed relative to the social optimum; however, they smoothly decline over time rather than exhibiting the sharp deadline distortion of the primary term.

7.2 Multiple wells per unit

Finally, we turn to the extended model in which each unit can accommodate up to ten wells. Figure 15 plots the mineral owner's value capture against the primary term length when units accommodate ten wells, holding the royalty fixed at 60%. In contrast to the single-well unit case (figure 11, panel (a)), the optimal primary term for all three initial price scenarios is infinite.

Why do primary terms perform so much more poorly for the mineral owner when units accommodate multiple wells? In this case, the net cost of drilling a well at the deadline is

Figure 15: Expected value to the mineral owner as a function of the primary term, with 10-well drilling units

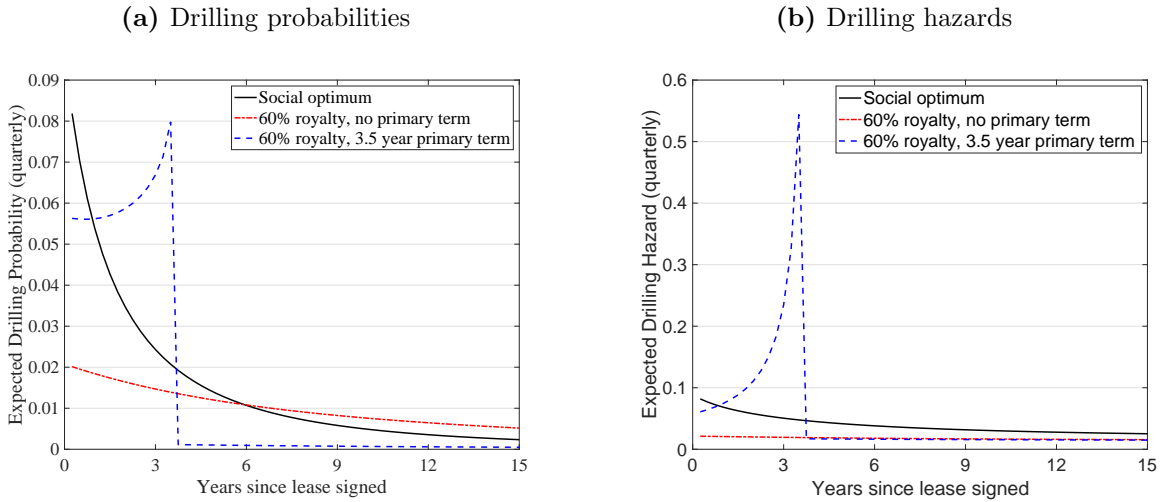


Note: Over all primary terms, the royalty is held fixed at 60%. The drilling subsidy paid by the mineral owner is \$0. All expected values are shown as a percentage of the expected value from the socially optimal drilling program. See text for details.

much lower than the bonus payment required to extend the lease, since the bonus is optimized to capture the value from drilling all of the potential wells, not just one of them. Thus, the primary term in this case will frequently bind on the first well, as shown in the simulated drilling probabilities in figure 16, panel (a) and especially in the simulated drilling hazard rates in panel (b). These spikes in drilling rates at the lease’s expiration date represent a large timing distortion that is especially costly in states of the world in which the firm would rarely choose to drill in the absence of the deadline.

Unlike the distortion that the primary term exerts on the first well to be drilled, it has no effect on the firm’s decision to drill the remaining wells on the lease. These wells therefore remain delayed in expectation, relative to the socially optimal drilling program, due to the royalty. Thus, when leases accommodate multiple wells, a primary term will tend to over-distort drilling of the initial well while failing to correct the royalty-induced distortions to

Figure 16: First-well drilling probabilities and hazards as a function of lease terms, when the lease accommodates 10 wells



Note: Drilling hazard each period is conditioned on not having drilled or let the lease expire prior to that period, as well as on accepting the initial lease terms (i.e., paying the bonus) in the first place. In all cases, the drilling subsidy paid by the mineral owner is \$0. See text for details.

later wells, thereby reducing the mineral owner’s returns in expectation.

In contrast, the drilling subsidy is still an effective tool for increasing the mineral owner’s expected revenue when units accommodate multiple wells. Because the subsidy applies to all wells, not just the first well, the impact of the subsidy on the mineral owner’s value capture does not change. Thus, figure 13, which presents owner value for varying drilling subsidies under the one-well case, continues to apply in the multi-well case.

Prior to the shale boom, spacing rules typically did not result in units that could accommodate such large numbers of non-interfering wells. However, due to the poor permeability of shale formations, standard one square mile spacing units will typically require several wells to fully extract the resources. Our result that primary terms substantially improve mineral owners’ expected revenue for one-well but not ten-well units suggests that contracting and spacing practices have not evolved in tandem with recent technological changes.

8 Conclusion

Oil and natural gas rights are typically owned by entities without the capital or expertise to extract and market the resources. As a result, the mineral rights owner contracts with a firm, leading to a principal-agent problem. In this paper, we describe the ubiquitous contract structure that has arisen in the United States and describe the distortions induced therein. We focus in particular on the primary term, which has been a purported driver of levels and location of oil and natural gas drilling activity. Using detailed data on private leases, pooling units, and well drilling, we demonstrate that expiration of the primary term clearly distorts the timing with which wells are drilled.

To rationalize the presence of these distortionary features, we model the mechanism design problem facing the mineral owner. We find that making the contract contingent on both realized production and the firm's drilling timing help the mineral owner maximize her revenue in the presence of asymmetric information and moral hazard. The revenue-optimal contract strikes a balance between extracting the firm's information rent and distorting its choice of drilling timing and inputs. The three parts of the contract are a tax on revenue, a drilling subsidy, and a lump-sum payment. Using the intuition of the linkage principle, the contract ties payments to observed production and timing, which are correlated with unobserved productivity and drilling effort.

Because drilling subsidies are likely precluded in practice due to mineral owners' liquidity constraints, we develop a computational model to examine how primary terms are a second-best tool for improving the owner's revenues relative to a royalty-only lease. We show that in situations that were common prior to the shale revolution, in which a single lease or pooling unit accommodated one well and no more, inclusion of a primary term shifts drilling forward in time so that it is closer to the socially optimal time path, albeit with a sharp discontinuity at the expiration date. This shift counteracts the royalty distortion and improves the expected revenue for the mineral owner while also increasing total surplus.

However, when drilling a single well allows a firm to retain an option to drill many future

wells—as is true in the Haynesville and other shale formations—primary terms no longer substantially improve the mineral owner’s expected payoff. This result suggests that mineral contracting has not kept pace with changes to the economic and institutional environment precipitated by the shale boom and is consistent with recent litigation brought by mineral owners regarding large pooling units that allow firms to hold large areas with a single well.³²

Our mechanism design result applies to a variety of settings beyond the oil and gas sector. For example, franchising agreements and intellectual property licenses often take a form reminiscent of the bonus-royalty-primary term contracts observed in our context. These assets exhibit similar properties to those driving our theoretical results. In the case of franchising, the franchisee may have better information about local demand conditions and can exert hard-to-observe but costly effort, while the franchisor can observe revenues and the time at which certain benchmarks (such as branch openings) occur. Licensors and licensees of intellectual property face similar constraints and asymmetries. Licensors are uncertain about the value of their IP, and licensees may prefer to wait to develop products utilizing acquired IP.

Looking ahead to future work, we first aim to improve our model’s calibration and use the model’s simulated drilling probabilities to better estimate the important parameters governing firms’ drilling decisions. In particular, we would like to incorporate endogenous fracking “effort”, measured with data on water injection, into the model so that we may simulate the effects of royalties on not just drilling timing but also input choices and output per well. We also hope to expand our geographic scope to include a larger share of the natural gas extraction industry, which would allow us to estimate the effects of primary term-induced drilling on overall U.S. natural gas supply and equilibrium natural gas prices. Finally, there is significant interest in understanding the supply elasticity of unconventional hydrocarbon resources (e.g., (Newell et al. 2016)). Aside from predicting future prices and quantities, this elasticity is crucial for understanding the effect of a number of policies such

³²*Gatti et al. vs State of Louisiana*, January 15, 2014.

as carbon taxes, methane regulations, and severance taxes. Our results suggest that the supply elasticity of a play may vary significantly over time due to dynamic contract terms, independently of costs or technological change.

References

- AGERTON, M. (2018): “Learning Where to Drill: Drilling Decisions and Geological Quality in the Haynesville Shale,” Working Paper.
- ANDERSON, S. T., R. KELLOGG, AND S. W. SALANT (2018): “Hotelling Under Pressure,” *Journal of Political Economy*, 126, 984–1026.
- BERMAN, A. (2015): “Haynesville Shale Needs \$6.50 Gas To Break Even: The Business Model Is Broken,” *Forbes*.
- BHATTACHARYA, V., A. ORDIN, AND J. W. ROBERTS (2018): “Bidding and Drilling Under Uncertainty: An Empirical Analysis of Contingent Payment Auctions,” Working Paper.
- BOARD, S. (2007): “Selling Options,” *Journal of Economic Theory*, 324–340.
- BOGAN, J. (2009): “Boom Times at the Haynesville Shale,” *Forbes*.
- BROWN, J., T. FITZGERALD, AND J. G. WEBER (2016): “Capturing Rents from Natural Resource Abundance: Private Royalties from U.S. Onshore Oil & Gas Production,” *Resource and Energy Economics*, 46, 23–38.
- CAMERON, A. C. AND P. K. TRIVEDI (2005): *Microeconometrics: Methods and Applications*, Cambridge University Press.
- CHETTY, R., J. N. FRIEDMAN, T. OLSEN, AND L. PISTAFERRI (2011): “Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records,” *Quarterly Journal of Economics*, 126, 749–804.
- CONG, L. W. (2018): “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, 936–959.

- GÜLEN, G., S. IKONNIKOVA, J. BROWNING, K. SMYE, AND S. W. TINKER (2015): “Production Scenarios for the Haynesville Shale Play,” *SPE Economics & Management*, 7.
- HAILE, P., K. HENDRICKS, AND R. PORTER (2010): “Recent U.S. Offshore Oil and Gas Lease Bidding: A Progress Report,” *International Journal of Industrial Organization*, 28, 390–396.
- HENDRICKS, K. AND D. KOVENOCK (1989): “Asymmetric Information, Information Externalities, and Efficiency: the Case of Oil Exploration,” *RAND Journal of Economics*, 20, 164–182.
- HENDRICKS, K. AND R. H. PORTER (1996): “The Timing and Incidence of Exploratory Drilling on Offshore Wildcat Tracts,” *American Economic Review*, 86, 388–407.
- HENDRICKS, K., R. H. PORTER, AND G. TAN (1993): “Optimal Selling Strategies for Oil and Gas Leases with an Informed Buyer,” *American Economic Review Papers and Proceedings*, 83, 234–239.
- HILLER, J. (2012): “Eagle Ford Lease Deadlines Driving Drilling,” *San Antonio Express-News*, San Antonio Express-News.
- KAISER, M. J. (2012): “Profitability assessment of Haynesville shale gas wells,” *Energy*, 38, 315–330.
- KAISER, M. J. AND Y. YU (2011): “Haynesville shale gas well performance and development potential,” *Natural Resources Research*, 20, 217–229.
- KALNINS, A. (2005): “Overestimation and Venture Survival: an Empirical Analysis of Development Commitments in International Master Franchising Ventures,” *Journal of Economics and Management Strategy*, 14, 933–953.

- KAPLOW, L. AND S. SHAVELL (2002): “On the Superiority of Corrective Taxes to Quantity Regulation,” *American Law and Economics Review*, 4, 1698–1734.
- KELLOGG, R. (2014): “The Effect of Uncertainty on Investment: Evidence from Texas Oil Drilling,” *American Economic Review*, 104, 1698–1734.
- LAFFONT, J.-J. AND J. TIROLE (1986): “Using Cost Observation to Regulate Firms,” *Journal of Political Economy*, 94, 614–641.
- (1987): “Auctioning Incentive Contracts,” *Journal of Political Economy*, 95, 921–937.
- LIN, C. C.-Y. (2013): “Strategic Decision-Making with Information and Extraction Externalities: a Structural Model of the Multistage Investment Timing Game in Offshore Petroleum Production,” *Review of Economics and Statistics*, 95, 1601–1621.
- MALE, F., A. W. ISLAM, T. W. PATZEK, S. IKONNIKOVA, J. BROWNING, AND M. P. MARDER (2015): “Analysis of gas production from hydraulically fractured wells in the Haynesville Shale using scaling methods,” *Journal of Unconventional Oil and Gas Resources*, 10, 11–17.
- MCAFEE, R. P. AND J. MCMILLAN (1987): “Competition for Agency Contracts,” *RAND Journal of Economics*, 18, 296–307.
- METCALF, G. E. (2010): “Investment in Energy Infrastructure and the Tax Code,” in *Tax Policy and the Economy, Volume 24*, ed. by J. R. Brown, University of Chicago Press.
- MILGROM, P. R. AND I. SEGAL (2002): “Envelope Theorems for Arbitrary Choice Sets,” *Econometrica*, 70, 583–601.
- MILGROM, P. R. AND R. J. WEBER (1982): “A Theory of Auctions and Competitive Bidding,” *Econometrica*, 50, 1089–1122.

- NEWELL, R. G., B. C. PREST, AND A. VISSING (2016): “Trophy Hunting vs. Manufacturing Energy: The Price-Responsiveness of Shale Gas,” NBER Working Paper No. 22532.
- PATZEK, T. W., F. MALE, AND M. MARDER (2013): “Gas Production in the Barnett Shale Obeys a Simple Scaling Theory,” *PNAS*, 110, 19731–19736.
- PORTER, R. H. (1995): “The Role of Information in U.S. Offshore Oil and Gas Lease Auctions,” *Econometrica*, 63, 1–27.
- RILEY, J. G. (1988): “Ex Post Information in Auctions,” *Review of Economic Studies*, 55, 409–429.
- ROBINSON, P. M. (1988): “Root-n-consistent Semiparametric Regression,” *Econometrica*, 56, 931–954.
- SKRZYPACZ, A. (2013): “Auctions with Contingent Payments—an Overview,” *International Journal of Industrial Organization*, 31, 666–675.
- WEITZMAN, M. L. (1974): “Prices vs. Quantities,” *Review of Economic Studies*, 41, 477–491.

Appendix

A Further Institutional Details

In Louisiana, pooling units are formation-specific. This means that the square-mile section Haynesville pooling unit only refers to the Haynesville. There may be other units in the same area with drilling targeting other formations. One complication is that a well that is drilled targeting one formation may lead the firm to have the rights to drill additional wells for the same lease that target different formations. This is due to the fact that in Louisiana, pooling units are formation-specific but leases typically (but not always) include all minerals under the ground, from any formation.

The fact that a lease can span multiple units vertically (through draining multiple formations) or horizontally (because a mineral owner leased land within multiple units as part of a single lease) implies that a firm can drill in one unit and hold production in another unit. Some leases include clauses known as Pugh clauses that restrict the ability of the firm to hold production in other units. For example, a vertical Pugh clause may restrict that the firm can only retain rights to the formation which it has drilled, leaving the mineral owner free to sign another lease with another firm to extract from other formations. In the case of a lease that spans multiple sections, a horizontal Pugh clause would restrict a firm that drills a Haynesville well on one section to have rights to drill on the other sections within the lease.

Conversations with DrillingInfo suggest that Pugh clauses are very rare. Firms typically approach mineral owners with a pre-written lease. Because it is not in the firm's interest to include a Pugh clauses, the default language typically excludes Pugh clauses. A mineral owner must be sufficiently savvy to negotiate a Pugh clause. It is very costly for us to check whether leases in our data have Pugh clauses: DrillingInfo's digitization of lease records does not include documentation of Pugh clauses, and therefore getting Pugh clauses would require us to visit each Louisiana parish to examine leases. Horizontal Pugh clauses (or the lack thereof) are unlikely to be important for this Haynesville analysis because few leases spanned multiple units/sections. Concerns about the lack of vertical Pugh clauses are mitigated because our sample excludes sections where there is production for non-Haynesville wells prior to the Haynesville boom. We also exclude sections where there is non-Haynesville drilling after the start of the Haynesville boom.

If there are multiple firms leasing land within a unit, the operating firm has a different role than the other firms. The operating firm is basically the decision maker. The remaining firms are non-operating participating shareholders, and share in both the revenues as well as in paying the costs of the firm. Non-operating participating shareholders have an incentive to share in the costs upfront, because failure to help pay for cost up front means that the operator can deduct 300% of the non-cooperative firm's share of cost before the operator disseminates to the non-cooperative firm its share of revenues.

Another type of non-operating participant are mineral owners within a unit that are not currently leasing their mineral rights. In contrast to firms leasing land on a unit that have an incentive to share in the costs, these mineral owners are not required to help pay for costs up front. Instead, the operating firm pays the unleased mineral owner her share of revenues

minus any costs.

B Data Appendix

In this data appendix, we first discuss our data sources. Then we discuss how we clean lease data. Finally we discuss how we match wells to PLSS sections.

B.1 Data sources

We take data from a number of sources, including the following:

- Louisiana DNR SONRIS data on well drilling and completions. We use publicly available data published by the Louisiana Department of Natural Resources (DNR). The well-level data include spud date, completion date, well name, formation targeted, top hole location, bottom hole location, and lateral location shapefiles. DNR SONRIS includes a data file that lists Haynesville wells. DNR SONRIS also includes a shapefile for PLSS sections.
- DrillingInfo production data. DrillingInfo takes unit-level reported monthly production data from the Louisiana DNR and then imputes well-level production using the start date of each well’s production.
- DrillingInfo lease data. DrillingInfo collects data on leases signed in Louisiana. Further details are below.
- Drilling cost and fracking input data. We obtain drilling cost information from reports (“Applications for Well Status Determination”) that unit operators file with the Louisiana DNR for the purpose of determining severance taxes. We obtain data on fracking inputs (water use and the number of frac stages used) from well completion reports that are available from the DNR.
- Dayrate data. We purchased these data from RigData, a U.S. onshore oil and gas industry intelligence service. We use dayrates that correspond to the “ArkLaTx” region, for rigs with depth ratings between 10,000 and 12,999 feet (which corresponds to the depth of the Haynesville).
- Henry Hub natural gas futures prices. We obtained daily futures price data, at all available delivery dates, at from Bloomberg.

B.2 Lease Data

We describe our steps to clean lease data. We discuss how we identify irrelevant observations as well as how we treat duplicates and likely duplicates.

The raw data of oil and gas leases in Louisiana was downloaded from DrillingInfo classic. We keep only leases in Bienville, Bossier, Caddo, De Soto, Natchitoches, Red River, Sabine, and Webster parishes—the parishes that cover the Louisiana portion of the Haynesville

formation. Data were downloaded in January 2016. Because we map leases to units which are square mile sections, we keep only those observations that report public land survey system township, range and section.

We keep observations that are listed as being leases, memo of leases, lease options, lease extensions, and lease amendments. We drop observations that are mineral rights assignments, lease ratifications, mineral deeds, royalty deeds, and other documents.

We drop leases with zero or missing acreage. Figure 17 shows the fraction of observations in our data with zero or missing acreage. Because leases with zero or missing acreage tend to be at the beginning of the Haynesville leasing period, we may be undercounting leasing at the beginning of the period. In that case, we may want to over-weight the earlier years of leases to compensate for the low rate




Figure 17: Raw lease data, showing the fraction of observations with zero or missing acreage and total number of leases.

Leases include information on the grantor of the lease (typically the original mineral owner) and the grantee (the oil and gas firm that leases the land). In some cases we find that oil and gas firms are listed as grantors, with other oil and gas firms listed as grantees. As these observations are likely cases where the land was re-leased or subleased, we drop these observations.

We also drop some other observations: We drop all excess observations that are perfect duplicates. We drop observations where information is inconsistent, such as where the reported township, range, and section is not within the stated reported parish. We drop leases that have lengths of less than 10 days.

We face some challenges with leases that span multiple sections. The data set is at the lease by section level. We identify leases that span multiple sections by identifying observations that share the same parish and record number. For those cases, we find that reported acreage is identical for all observations sharing parish and record number. This

implies that the data is reporting total acreage over the entire lease, rather than the total acreage of the lease within a given section. Therefore to impute within-section acreage of the lease, we assume that within-section acreage is proportional to total section area, and divide the total acreage by the number of sections to get at an updated area of the lease within each reported section.

We also examine other cases where a lease likely spans multiple sections but does not have a common parish or record number. This may happen if a lease spans multiple parishes or record number was mis-typed. To identify such cases, we identify observations that share both the same grantor name and the same grantee name, as well as have the same reported acreage. We spot check to make sure these cases are close together geographically. As above, we impute within-section acreage of the lease by assuming that within-section acreage is proportional to total section area.

We also find that in some cases, a single firm grantee has leased from multiple grantors and the reported acreage appears to be the total over all grantors. We identify those leases by identifying duplicates that share the same grantee name and the same acreage, and where the acreage reported is unusual—e.g., is large and/or is not equal to a multiple of common lot sizes. Similar to our approach described in the previous paragraph, here we impute a new acreage measure by dividing the reported acreage by the number of apparent duplicates.

After doing this, we find that total leased acreage still sometimes adds up to more than the total acreage of the section, and sometimes significantly so. One reason for this is that there may be separate observations for cases where for a given plot, there are multiple grantors (e.g., husband and wife), and separate observations are recorded for each grantor. Another reason for this is that in some cases it appears that data was entered multiple times and inconsistencies were never reconciled, such that two observations may differ only slightly.

To identify these likely duplicates, we use a hierarchical clustering method described by <https://stat.ethz.ch/R-manual/R-devel/library/stats/html/hclust.html>. In particular, we use the `hclust` function within the `cluster` package, version 2.0.7-1, for R. The `hclust` function uses information on how similar multiple observations are to each other to determine whether they are likely duplicates. The algorithm puts observations that are likely duplicates into the same “cluster”; from there we use proportional downweighting of all observations within the same cluster to get updated acreage. This method relies on constructing some kind of measure of similarity between any two observations i and j . Depending on the threshold level of similarity that the researcher imposes, the number of clusters can range from the total number of observations (no clustering) to 1 (all observations are placed within the same cluster).

This similarity measure we use is a Euclidean-like distance measure where the distance between observation i and observation j takes the form:

$$d_{ij} = \sqrt{\sum_k w_k m_k(x_i^k, x_j^k)} \quad (5)$$

Here k indexes characteristics of the observation—e.g., grantor name, the start date of the lease, the acreage, the reported royalty rate, etc. The function m_k is a function that determines how similar two observations are, and is equal to 0 if identical, and positive otherwise. Depending on the characteristic, we use different types of m_k functions:

- $m_k(x_i^k, x_j^k) = (x_i^k - x_j^k)^2$ for some numerical characteristics like the start date of the lease. Prior to inputting variables x^k into this function, we standardize them so that they have a mean of zero and standard deviation of one.
- $m_k(x_i^k, x_j^k) = 1(x_i^k \neq x_j^k)$ for other numerical and binary characteristics like reported royalty rate, acreage, and whether there is an extension option.
- $m_k(x_i^k, x_j^k)$ is a fuzzy match score for string characteristics like grantee name and grantor name. We use the `partial_ratio` function from the `fuzzywuzzy` Python package, version 0.16.0. The `partial_ratio` function uses Levenshtein distance augmented with partial string matching. It allows us to identify cases where some subsets of words within strings match or nearly match, even if the length of the two strings is very different. This is useful for catching cases with identical last names but differing first names. We scale this measure so that it ranges from 0 to 1.

For cases where information is missing, we set a value of $m_k = 0.4$ if both observations are missing and $m_k = 0.7$ if only one observation is missing.

w_k are positive weights. In our base specification, we set $w_k = 1$ for all characteristics other than acreage, for which we set $w_k = 100$. This ensures that leases that vary in acreage will not be presumed to be duplicates.

Using this weighting method, we find in spot checks a few cases where observations likely should have been combined in the same cluster but were not because they had substantially different grantor names. Such could happen if two grantors are siblings where through marriage a last name changes. In such cases, the `partial_ratio` would give a very high score. This implies that we should use a lower w_k weight for grantor names. We also find very little variation in grantee name within section, which suggests leases that do differ in grantee name are very unlikely to be duplicates, suggesting that the weight w_k for grantee names should be higher. Therefore for robustness, we do an alternative weighting where we set $w_k = 1/2$ for grantor name and $w_k = 2$ for grantee name, with all other weights the same. Results using this approach are very similar.

How many observations are clustered together depends on the threshold level of similarity imposed by the researcher. Figure 18 is an example from one section that shows that the number of clusters depends on the threshold similarity level (shown on the y axis), which we refer to as the threshold height. Figure 19 shows a similar graph for a more heavily leased section.

Therefore, we need to determine what the “right” threshold height is, which we refer to as our preferred threshold height. To do so, we choose a calibration date of January 1, 2010, examining only the leases that were active on that date.³³ We first examine every possible threshold height that could be used to cluster the leases in the section. For each possible threshold height, we find the resulting clusters, and then downweight each lease’s acreage by the total number of observations in the cluster, and then compute what total acreage would be within the section. Then, for each section, we find the threshold height would be that

³³We find that using other calibration dates gives similar results. We use January 1, 2010, as it was at a period of peak leasing, and therefore a period where it is most likely that most of a section had been leased. Leases whose primary terms would have expired but may have been extended are not included in this group.




Figure 18: Dendrogram visualization of clustering of a section with few observations during July 1, 2009. Different colors represent different clusters. The dotted line is our preferred threshold height that we describe below.




Figure 19: Dendrogram visualization of clustering of a more heavily leased section leased during July 1, 2009. Different colors represent different clusters. The dotted line is our preferred threshold height that we describe below.

would set total acreage leased to be equal to or just under total section area. We refer to this as the section-level threshold height. Then we set our preferred threshold height to be equal to the 90th percentile of all the section-level threshold heights.³⁴

Our preferred threshold height is 1.65. Given that most of our weights are 1, this implies that if two leases only differ in one categorical characteristic, they will be combined into one cluster. But if they differ in two or more, they will not be combined.

From here we apply our preferred threshold height for a series of dates in our sample: We examine quarterly dates ranging from January 1, 2005 to January 1, 2016. For each date, we do agglomerative clustering using our preferred threshold height and output the set of clusters. This gives us lease by date-specific downweights. We find that in some cases lease downweights vary depending on the date: For example, a lease may be in a cluster of five on April 1 2010 but a cluster of six July 1 2010—resulting in a downweight of 1/5 for April 1 2010 and a downweight of 1/6 for July 1 2010. In those situations, we take the inverse of the arithmetic average of the inverse over all quarterly dates to get a master downweight for each lease—e.g., to arrive at a weight of 1/5.5.

In some outlier cases we find that even with this downweighting, total implied acreage still exceeds section acreage. This happens because for some sections, we find large individual reported acreage for multiple leases. For those cases we further downweight area using the implied total leased area so that total acreage is equal to section acreage in the most heavily leased date, and less for other dates.

Finally we use our estimated downweight to update our data set: Our new measure of acreage is equal to the old measure multiplied by the downweight.

B.3 Wells and Well-Unit Matching

Here we discuss how we construct well data and match wells to PLSS sections. Because in the Haynesville formation, unit boundaries are the PLSS section boundaries, matching wells to units is analogous to matching wells to PLSS sections.

We differentiate between wells targeting the Haynesville formation and wells targeting other formations. This is because units for the Haynesville Shale were almost always one square mile sections, whereas units for the other formations may have had other boundaries. We identify Haynesville wells from three sources: The first is an auxiliary DNR file that is limited Haynesville wells. The second and third are from the main well data. The second is the name of the well: wells targeting the Haynesville typically have names that begin with “HA” or “HAY”. The third is whether the listed formation that the well targets. We identify a well as targeting the Haynesville if it satisfies at least one of these criteria. We also restrict that Haynesville wells must have been spudded on or after September 2006.

To match wells to sections, we use information on the reported laterals, reported bottom holes, and reported top hole location. If for a given well, the data only reports top hole location, we use the location of the top hole to identify which PLSS section the well is in. If the data reports bottom hole but not lateral information, we use the location of the bottom hole to identify which PLSS section the well is in. If the data reports lateral information,

³⁴We find very similar results if we use a threshold height of 85%. We find that using the 100th percentile is not feasible as some outlier sections have leases with large acreage such that total leased acreage is always greater than section acreage unless leases of different acreage are combined into the same clusters.

we use the PLSS section that the lateral runs through to identify the PLSS section of the well. Because the vast majority of our data has reported laterals, our results are robust to other approaches.

In some cases, the PLSS section intersects multiple sections. This may be due to two different reasons. One is that the well top hole is located in a different section than the section the well extracts from, and the top hole and initial lateral begin in a different section simply for the purpose of sharing a well pad with other wells or to give sufficient space to accommodate the curvature of transitioning from the vertical to the horizontal while still extracting from a maximum area within the targeted PLSS section. A second reason is that the well actually targets multiple PLSS sections. Therefore in those cases, we map the well to each PLSS section for which it at least 300 meters of the horizontal well bore passes through the PLSS section. Because there are very few wells that have significant acreage that passes through multiple PLSS sections, our results are robust to alternative definitions.

C Additional Empirical Analysis

In this Appendix, we discuss additional descriptive results using data from the Louisiana portion of the Haynesville shale.

C.1 Lessee and Operator Shares

Here we briefly discuss acreage shares for original lessees listed in the lease documents. We find that there are a large number of lessees, many leasing very small acreage. The right column of Table 6 gives shares for those lessees who have at least 2% of acreage. These shares are acreage based, where acreage is derived as discussed in Appendix subsection B.2.

We also examine the unit operators who eventually make the drilling decisions for these leases, as well as the mapping from lessee to operator. We find that unit operators are a bit more concentrated. Operator shares of these leases are listed in the last row of Table 6. We find that some of the largest operators, including Chesapeake and Petrohawk, are also major lessees. In contrast, Encana and SWEPI got much of their acreage from other players. Each column of Table 6 lists the fraction of operator acreage that came from each of the major lessees. Operators with shares less than 2% are not listed.

C.2 Bunching Analysis

To show that the drilling spike prior to primary term expiration is statistically significant, we use a bunching estimate similar to that of Chetty et al. (2011). We take time of spud relative to first lease expiration date, discretize it to the quarterly level, and compute total wells spudded for each quarter. This results in a total sample size of 34 quarters. We create some indicator variables for whether the spud date is two quarters before lease expiration (`pre_2`), one quarter before lease expiration (`pre_1`), one quarter after lease expiration (`post_1`), and two quarters after lease expiration (`post_2`). We also add similar variables for spud timing relative to the extension expiration date (`pre_ext2`, `pre_ext1`, `post_ext1`, and

Lessees:	Operators:						Lessee share
	Chesapeake	Petrohawk	Encana	Samson	SWEPI	Others	
Chesapeake	0.375	0.078	0.030	0.025	0.137	0.102	0.179
Petrohawk	0.054	0.371	0.034	0.075	0.020	0.060	0.107
Pride	0.005	0.044	0.400	0.038	0.009	0.030	0.068
Audubon	0.007	0.002	0.104	0.001	0.257	0.002	0.028
Honeycutt	0.048	0.021	0.005	0.009	0.098	0.008	0.028
Suncoast	0.047	0.011	0.000	0.000	0.000	0.021	0.023
SGS	0.061	0.004	0.004	0.000	0.000	0.003	0.023
Encana	0.000	0.011	0.141	0.003	0.024	0.008	0.022
Samson	0.001	0.004	0.002	0.261	0.002	0.006	0.022
Long	0.029	0.007	0.002	0.000	0.014	0.036	0.021
Others	0.373	0.448	0.279	0.589	0.439	0.724	0.480
Total	1	1	1	1	1	1	
Oper. share	0.342	0.172	0.120	0.071	0.047	0.247	

Table 6: Lessee and operator shares, where shares are acreage based. The central cells give the fraction of a given operator’s acreage that comes from a given lessee. Bolded entries denote where the lessee and the operator are the same. The right column is lessee shares. The bottom row is operator shares. Shares are only listed if the lessee or operator has more than 2% shares. HHI for lessees is 0.056 and HHI for operators is 0.172. Shares are acreage based, where acreage is derived as discussed in Appendix subsection B.2.

post_ext2), assuming that the extension expiration date is 2 years after the primary term expires.

Then we estimate a regression of the form:

$$c_t = f(t) + \beta_1 \cdot \text{pre}_2 + \beta_2 \cdot \text{pre}_1 + \beta_3 \cdot \text{post}_1 + \beta_4 \cdot \text{post}_2 + \beta_5 \cdot \text{pre_ext2} + \beta_6 \cdot \text{pre_ext1} + \beta_7 \cdot \text{post_ext1} + \beta_8 \cdot \text{post_ext2} + \varepsilon_t \quad (6)$$

where c_t is total well count, t is quarter, and $f(t)$ is a polynomial of degree 9. Our regression estimates are in Table 7. The estimates of β_1 , β_2 , β_5 , and β_6 are all statistically significant with p values less than 0.05, showing that there is significantly more drilling in the two quarters prior to the primary term expiration and the two quarters prior to any extension term expiration. We also do a number of joint tests, including tests that β_1 through β_4 are all jointly equal to zero, whether β_5 through β_8 are all jointly equal to zero, and whether β_1 through β_8 are all jointly equal to zero. For each test, we get p values that are less than 10^{-30} . In Figure 20, we plot our data as well as our polynomial predicted probabilities.

pre_2	14.25 (5.37)
pre_1	67.04 (5.71)
post_1	8.69 (5.61)
post_2	-6.80 (5.11)
pre_ext2	13.30 (5.42)
pre_ext1	18.99 (5.86)
post_ext1	-4.96 (5.72)
post_ext2	-3.73 (5.00)
R squared	
Observations	35

Table 7: Bunching regression estimates. Standard errors are in parentheses.

C.3 Information Externalities and Common Pools

One alternative explanation for the drilling patterns we observe is some kind of externality, such as information externalities or common pools. Common pools and information externalities shift the predicted drilling patterns in different ways: If common pools exist, we might expect a race to drill. If information externalities are driving results, we might expect reluctance to drill, as firms wait for nearby firms to first drill before drilling.

To examine the effect of externalities, we compare units where most of the unit’s neighbors are operated by the same operator with units where most of the unit’s operators are operated by a different operator. We would expect that the latter are more likely to have these externalities present. We define neighboring units as those for which the centroids of the two units are within 1.2 miles. (Results are very similar when using a threshold of 1.7, which will include the diagonal units.)

In Figure 21, we compare drilling probabilities for these two types of units. We find that units where neighboring operators tend to be different have slightly higher probabilities of

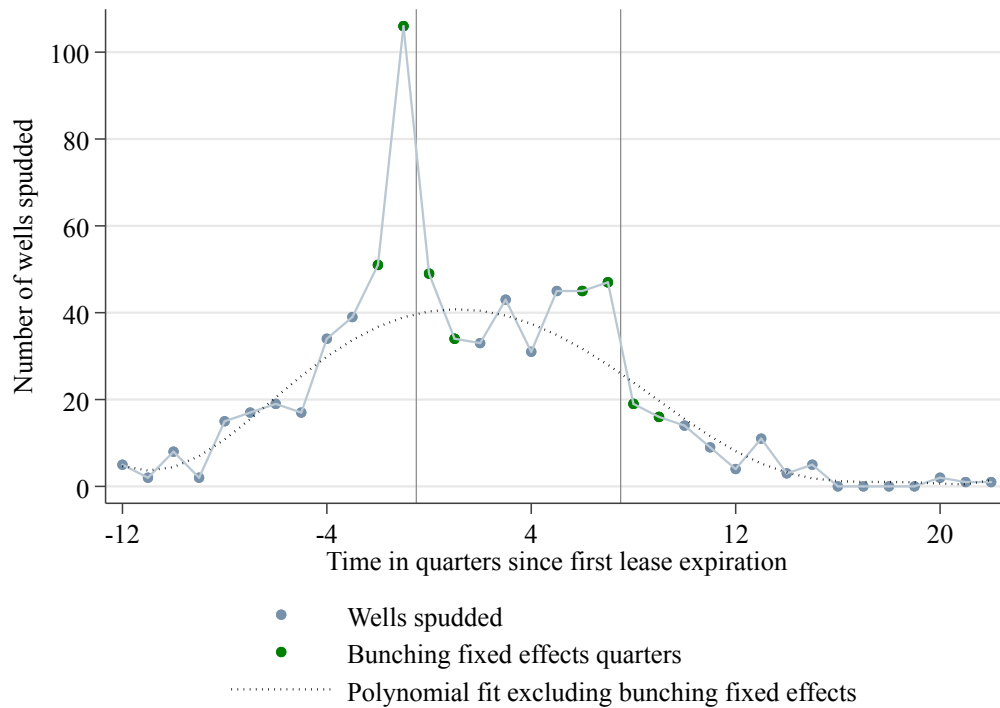


Figure 20: Figure demonstrating our bunching analysis, including wells spudded count data, which quarters we add bunching fixed effects, and the polynomial predicted probabilities given bunching estimator fixed effects. Timing is relative to the expiration date of the first lease within the unit to expire. Vertical lines are drawn at the date of first lease expiration and two years after first lease expiration.

drilling in earlier months, suggesting that the common pool mechanism is dominating the information externality mechanism. However for both groups, we find a sharp rise in the probability of drilling in the few months prior to the primary term expiring. Therefore we conclude the externalities are not driving these results.

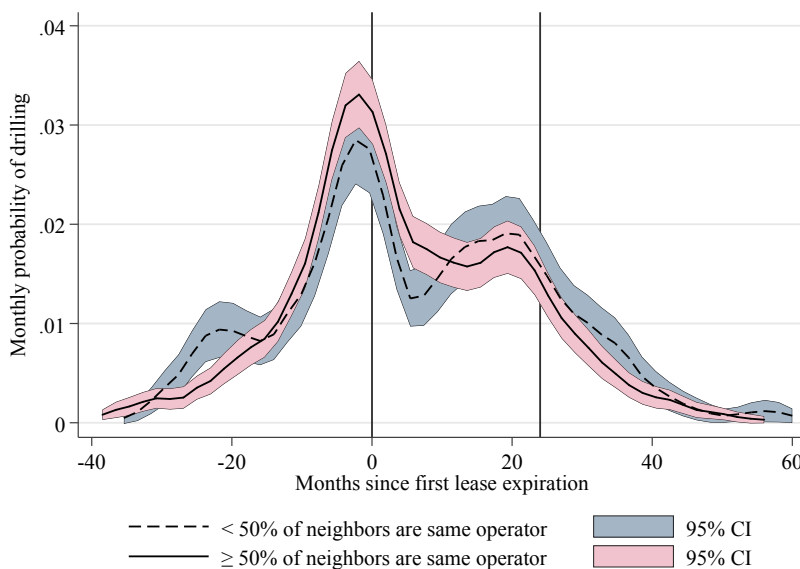


Figure 21: Smoothed probability estimates comparing units where most of the nearby units have the same operator vs. units where most of the nearby units have a different operator. Neighboring units are defined as those with centroids within 1.2 miles of the centroid of the given unit. Graph shows the date of the first Haynesville well drilled within a section relative to the expiration date of the first lease within the unit to expire. Vertical lines are drawn at the date of first lease expiration and two years after first lease expiration.

C.4 Drilling Inputs and Outputs as a Function of Drilling Time

Examine the values of well inputs (water used in fracking and total well costs) as well as total production. In particular, we examine how these variables change depending on whether drilling is done early, just prior to the first primary term expiring, or after the first primary term expires. Overall we find that these variables are very noisy, and for the most part do not vary significantly depending on the time of drilling.

Figure 22 graphs the first 12 months of production on the y axis, where the x axis is the time of drilling relative to the date that the first primary term expires. We find that average productivity declines somewhat as drilling time approaches the primary term, but this appears to be largely driven by outlier wells with low productivity. Wells drilled into the extension period have somewhat higher production.

Figure 23 graphs water use for hydraulic fracturing. We find that log water use gradually decreases with time of drilling relative to primary term for cases where the well is drilled prior to primary term construction. This is consistent with the hypothesis a firm that




Figure 22: Graph plots $\log(\text{Mcf gas production})$ over the first 12 months of the life of a section’s first well. The x-axis is measured in days relative to the first lease expiration at the time the well was spudded. The unit of analysis is the section. The line is the predicted value from a local polynomial regression.

is drilling only to hold the lease by production may use less inputs. We see a similar period during the last year of the extension period (days 365 to 730 in the graph) with declining water inputs.

Figure 24 graphs reported well drilling (accounting) costs.³⁵ Reported well costs decline as a function of the lease time, but are quite flat for the final year or so of the primary term.

C.5 Distribution of the Expiration Dates of Leases

All units have multiple leases on them. Typically leases are signed at different dates and therefore the primary term tend to end on different dates. Figures 5, 6, 7, 9 and 10 examine the timing of drilling relative to the number of months since the first well started. In this subsection, we show that the distribution of the expiration dates has an effect on the timing of drilling in the model.

For each unit, we calculate the amount of time from when the first lease expires to when the median primary term expiration date, where the median is calculated on an acreage weighted basis. We compute that median of that statistic. From there, we compare drilling timing on leases that had a relatively long time until median expiration (time is above the median) with those that had a relatively short time until median expiration (time is below the median). As shown in 25, we find that leases that had a longer time until median

³⁵These reported costs come from the DNR and are submitted by firms as part of their applications for severance tax exemptions. These documents are known as “Applications for Well Type Determination” documents. We downloaded these from DNR’s SONRIS website.




Figure 23: Graph plots $\log(\text{gallons of water})$ used in fracking a section's first well. The x-axis is measured in days relative to the first lease expiration at the time the well was spudded. The unit of analysis is the section. The line is the predicted value from a local polynomial regression.




Figure 24: Graph plots $\log(\text{reported drilling cost})$. The x-axis is measured in days relative to the first lease expiration at the time the well was spudded. The unit of analysis is the section. The line is the predicted value from a local polynomial regression.

exploration had less of a spike in drilling in the few months prior to the first lease expiring. This show that the distribution of lease expiration dates has an effect on drilling timing.




Figure 25: Smoothed probability estimates of drilling probability timing comparing sections where it took above the median amount of time to reach 50% of leases expiring relative to places where it took less than the median amount of time for 50% of leases to expire. Graph shows the date of the first Haynesville well drilled within a section relative to the expiration date of the first lease within the unit to expire. Vertical lines are drawn at the date of first lease expiration and two years after first lease expiration.

D Derivation of the optimal contingent contract for the principal

This appendix derives analytically the optimal contract for the principal (mineral owner), given the model set up from Section 4.1. The exposition below closely follows Board (2007), while including the elements of ex-post observable production and hidden effort that parallel Laffont and Tirole (1986). To facilitate the derivation of the optimal contract, we follow the standard approach of considering a direct revelation mechanism in which the firm reports a type $\hat{\theta}$ and is then assigned an up-front “bonus” transfer of $R(\hat{\theta})$ at $t = 0$ and a contingent payment $z_t(\hat{\theta}, y, P^t)$ to be paid when the option is executed. For now, we allow this payment to be contingent on the reported type, ex-post production, and the price history up to execution, though in practice conditioning only on the first two arguments will be necessary for optimality.

D.1 Firm’s problem

The firm must make three decisions, in sequence:

1. Report a type $\hat{\theta}$ to the principal at $t = 0$ (or opt-out)
2. Choose a time $\tau \in \{1, \dots, T\}$ at which to exercise the option to drill, where $\tau = \infty$ signifies not drilling.
3. Conditional on drilling, select an effort level $e \in \mathbb{R}^+$

Let $\tau^*(\theta, z)$ denote a decision rule that dictates whether the well should be drilled in each period t given the gas price P_t (we suppress the dependence of z on $\hat{\theta}$, y , and P^t). The firm's problem, conditional on participation, is then given by

$$\max_{\hat{\theta}, \tau^*(\theta, z), e} \Pi(\hat{\theta}, \tau^*(\theta, z), e, \theta) = E_P[(P_\tau g(e)\theta - (c_0 + c_1 e) - z_\tau(\hat{\theta}, g(e)\theta(1 + \varepsilon), P^t))\delta^\tau] - R(\hat{\theta}), \quad (7)$$

where E_P is the expectation at the start of period 1, taken over all prices. Note that social surplus is maximized by the solution to (7) when the transfers z and R are set to zero.

From the definition of $g(e)$ in Section 4.1, it is clear that the effort selection problem has a unique, interior solution. In addition, the decision rule $\tau^*(\theta, z)$ will be given by an optimal stopping rule.³⁶

As usual, we restrict attention to truth-telling mechanisms that induce the agent to report $\hat{\theta} = \theta$. Let $\tau(\hat{\theta})$ and $e(\hat{\theta})$ denote the timing rule and effort function that correspond to the optimal truthful mechanism. Because drilling is observable, $\tau(\hat{\theta})$ can be imposed by the principal. For truth-telling to be incentive compatible, it must be the case that $e(\hat{\theta})$ (which is not observed by the principal and therefore cannot be imposed) is the optimal effort level for the firm, subject to the mechanism.

To characterize the firm's ability to deviate from $e(\hat{\theta})$ and thereby reap information rent, we follow Laffont and Tirole (1986) by first restricting attention to deviations in a *concealment set* in which, for any report $\hat{\theta}$, the chosen effort \tilde{e} is such that $g(\tilde{e})\theta = g(e(\hat{\theta}))\hat{\theta}$. Thus, absent uncertainty generated by ε , any deviation outside the concealment set can be detected by the principal.³⁷

Within the concealment set, the firm's choice of report $\hat{\theta}$ determines the firm's effort level \tilde{e} . Define an inverse production function $H(E)$ by $g(H(E)) = E$. Note that the derivatives of g and H are related by $H'(g(e)) = 1/g'(e)$. The firm's problem may then be written:

$$\max_{\hat{\theta}} \Pi(\hat{\theta}, \theta) = E_P[(P_{\tau(\hat{\theta})} g(e(\hat{\theta}))\hat{\theta} - (c_0 + c_1 H(g(e(\hat{\theta})))\frac{\hat{\theta}}{\theta})) - z_{\tau(\hat{\theta})}(\hat{\theta}, g(e(\hat{\theta}))\hat{\theta}(1 + \varepsilon), P^{\tau(\hat{\theta})})\delta^{\tau(\hat{\theta})}] - R(\hat{\theta}). \quad (8)$$

To obtain the marginal information rent for a firm of type θ , we need the derivative of $\Pi(\hat{\theta}, \theta)$ with respect to θ . Using the envelope theorem, this marginal rent is given by the

³⁶Board (2007) proves the existence of such a rule for the case in which effort e is fixed. Existence in our model follows the same proof, with the assumptions that P_t and $g(e)$ are bounded above replacing the Board (2007) assumption that costs are bounded below.

³⁷In the presence of a non-degenerate distribution $\Lambda(\varepsilon)$, the sufficiency condition for implementing the mechanism will be stricter than the conditions given below that the optimal stopping time is decreasing in θ and that $\partial e(\theta)/\partial \theta \geq 0$ (also see footnote 40).

partial derivative:³⁸

$$\left. \frac{\partial \Pi(\hat{\theta}, \theta)}{\partial \theta} \right|_{\hat{\theta}=\theta} = E_P \left[\frac{c_1 g(e(\theta)) \delta^{\tau(\theta)}}{g'(e(\theta)) \theta} \right] \quad (9)$$

Equation (9) is the first-order incentive compatibility condition. The second order monotonicity condition is given by

$$\frac{\partial \Pi(\hat{\theta}, \theta)}{\partial \theta \partial \hat{\theta}} \geq 0. \quad (10)$$

From equation (9), it is apparent that the principal's mechanism will satisfy this condition if the optimal stopping time is decreasing in θ and if $\partial e(\theta)/\partial \theta \geq 0$.

To see that satisfaction of equation (10) is sufficient for incentive compatibility, consider a type θ firm that deviates to report type $\hat{\theta} \neq \theta$. We have:

$$\begin{aligned} \Pi(\hat{\theta}, \theta) &= \Pi(\hat{\theta}, \hat{\theta}) + \int_{\hat{\theta}}^{\theta} \frac{\partial \Pi(\hat{\theta}, s)}{\partial \theta} ds \\ &\leq \Pi(\hat{\theta}, \hat{\theta}) + \int_{\hat{\theta}}^{\theta} \frac{\partial \Pi(s, s)}{\partial \theta} ds \\ &= \Pi(\theta, \theta), \end{aligned} \quad (11)$$

where the second line follows from the monotonicity condition.

Given incentive compatibility, integration of equation 9 yields the firm's information rent:

$$\Pi(\theta, \theta) = E_P \left[\int_{\underline{\theta}}^{\theta} \frac{c_1 g(e(s)) \delta^{\tau(s)}}{g'(e(s)) s} ds \right], \quad (12)$$

where $\underline{\theta}$ denotes the lowest type that participates, so that $\Pi(\underline{\theta}, \underline{\theta}) = 0$.

D.2 Revenue-maximizing contract for the principal

Continuing to follow Laffont and Tirole (1986) and Board (2007), we treat the principal's problem as an optimal control problem in which the objective is to find $\tau(\hat{\theta})$ and $e(\hat{\theta})$ such that the expectation of social surplus minus information rent is maximized. We therefore write the principal's problem as

$$\max_{\tau(\theta), e(\theta, P), \underline{\theta}} \int_{\underline{\theta}}^{\bar{\theta}} \left[E_P (P_{\tau} g(e(\theta)) \theta - c_0 - c_1 e(\theta)) \delta^{\tau(\theta)} - \int_{\underline{\theta}}^{\theta} \frac{c_1 g(e(s)) \delta^{\tau(s)}}{g'(e(s)) s} ds \right] f(\theta) d\theta, \quad (13)$$

where the principal also chooses the type $\underline{\theta}$ for which the individual rationality constraint binds with equality.

To eliminate the double integral, we can use Fubini's theorem. Letting $h(\theta) \equiv f(\theta)/(1 -$

³⁸Following Board (2007), we use generalised envelope theorem of Milgrom and Segal (2002) because the space of stopping times is more complicated than \mathbb{R}^+ .

$F(\theta)$) denote the hazard function, we re-write the principal's problem as:

$$\max_{\tau(\theta), e(\theta, P), \underline{\theta}} \int_{\underline{\theta}}^{\bar{\theta}} E_P \left[\left(P_\tau g(e(\theta)) \theta - c_0 - c_1 e(\theta) - \frac{c_1}{\theta h(\theta)} \frac{g(e(\theta))}{g'(e(\theta))} \right) \delta^{\tau(\theta)} \right] f(\theta) d\theta. \quad (14)$$

At this point, it is useful to recall the firm's problem, equation (7). Following the logic in Board (2007), the principal can induce the firm to follow the stopping rule implied by (14) by setting the contingent payment z equal to the information rent term in (14), because doing so makes the firm's problem equivalent to the principal's problem. Specifically, the revenue-maximizing contingent payment is given by:³⁹

$$z_\tau(\theta, y, P^t) = \frac{c_1}{\theta h(\theta)} \frac{g(e(\theta))}{g'(e(\theta))}. \quad (15)$$

The optimal contingent payment z from the firm to the principal is positive, which will lead to delayed drilling relative to the social optimum. Note that the optimal payment is zero for the highest type $\bar{\theta}$ firm because $\frac{1}{h(\bar{\theta})} = 0$, reflecting the standard “no distortion at the top” rule. The optimal up-front payment $R(\theta)$ is set to equate the firm's utility to the information rent expressed in equation (12), where the utility of the endogenously chosen type $\underline{\theta}$ firm is set to zero.

The presence of a contingent payment upon execution of the option echoes Board's (2007) result. What differs here is that the optimal payment is contingent not just on drilling but also on effort. Indeed, holding type fixed, it is easy to verify that the payment in equation (15) is increasing in effort. Because effort is not observable to the principal, this optimal mechanism must be implemented using a payment that is contingent on production y rather than effort itself. Paralleling Laffont and Tirole (1986), we now examine an implementation that involves an affine contingent payment: a lump sum transfer combined with a linear tax on production. The appeal of an affine payment is that its optimality is robust to the distribution of the disturbance ε . The downside is that the sufficient conditions for incentive compatibility will be stronger than those discussed in Section D.1 above, because the affine payment structure constrains punishments for deviations outside of the concealment set.⁴⁰

To derive the optimal linear production tax, we first take the pointwise derivative of (14) with respect to $e(\theta)$ to obtain the FOC that defines the principal's optimal effort function, conditional on drilling at τ . Suppressing the dependence of $g(e(\theta))$ and its derivatives on θ ,

³⁹The principal's optimal timing and effort functions as defined by equations (14) and (16) will satisfy the firm's second order conditions (the optimal stopping time is decreasing in θ and $e'(\theta) \geq 0$) if the hazard rate is monotonically increasing ($h'(\theta) \geq 0$) and if $g''' \leq 0$. To see this sufficiency, observe that the total derivative of the term in parentheses in (14) is strictly increasing in θ , via application of the envelope theorem to the firm's problem. Thus, per Lemma 1 in Board (2007), the optimal stopping time is decreasing in θ . Second, take the derivative of (16) with respect to θ and collect terms involving $e'(\theta)$. It then becomes apparent that $h'(\theta) \geq 0$ and $g''' \leq 0$ are sufficient for $e'(\theta) > 0$.

⁴⁰The sufficient conditions for incentive compatibility of the affine contingent payment (equation (21)) are difficult to characterize in terms of primitives. In the event that they are not satisfied, the principal will need to “iron” over regions in the type space where incentive compatibility does not hold.

this FOC is given by:

$$FOC_{e(\theta)} : P_\tau g' \theta - c_1 - \frac{c_1}{\theta h(\theta)} \frac{g'^2 - gg''}{g'^2} = 0 \quad (16)$$

In FOC (16), the last term will be positive due to the assumption that $g'' < 0$, and the first two terms together will equal zero at the socially optimal effort level. Thus, for equation (16) to hold, $e(\theta)$ must be strictly less than socially optimal effort except for type $\bar{\theta}$, where the fact that $\frac{1}{h(\bar{\theta})} = 0$ implies that the highest type firm will exert optimal effort. Again, this is the standard “no distortion at the top” result.

To obtain the optimal production tax, we return to the firm’s problem, taking the derivative of equation (7) with respect to e to derive the firm’s FOC for its optimal effort, conditional on drilling at τ :

$$FOC_e : P_\tau g' \theta - c_1 - \frac{\partial z_\tau(\theta, y, P^t)}{\partial y} g' \theta = 0 \quad (17)$$

Combining equations (16) and (17) yields the linear tax on production that aligns the firm’s incentives with the effort function that the principal wishes to induce:

$$\begin{aligned} \frac{\partial z_\tau(\theta, y, P^t)}{\partial y} &= \frac{c_1}{\theta^2 h(\theta) g'} \frac{g'^2 - gg''}{g'^2} \\ &= \frac{c_1}{\theta^2 h(\theta) g'} \left(1 - \frac{gg''}{g'^2}\right) \end{aligned} \quad (18)$$

We can now solve for the optimal affine contingent payment using equations (15) and (18):

$$z_\tau(\theta, y, P^t) = \frac{c_1}{\theta h(\theta)} \frac{g}{g'} + \frac{c_1}{\theta^2 h(\theta) g'} \left(1 - \frac{gg''}{g'^2}\right) (y - g\theta) \quad (19)$$

Rearranging, we obtain:

$$z_\tau(\theta, P^t, y) = \frac{c_1 g}{\theta h(\theta) g'} \frac{gg''}{g'^2} + \frac{c_1 \left(1 - \frac{gg''}{g'^2}\right)}{\theta^2 h(\theta) g'} y \quad (20)$$

Note that in equation (20), both terms of the contingent payment are zero for the highest type, again reflecting the usual “no distortion at the top” result.

Finally, we may use equation (17) to eliminate c_1 in equation (20) and obtain:

$$z_\tau(\theta, P^t, y) = \frac{gg'' P_\tau g \theta}{\theta h(\theta) g'^2 + g'^2 - gg''} + \frac{(g'^2 - gg'') P_\tau y}{\theta h(\theta) g'^2 + g'^2 - gg''} \quad (21)$$

The second term in (21) is a positive tax on revenue $P_\tau y$; i.e., a royalty. The first term is negative (because $g'' < 0$) and represents a payment from the principal to the firm at the time of drilling that is not dependent on output. The affine implementation of the optimal mechanism therefore taxes natural gas production and subsidizes drilling, where the tax and subsidy rates are type-dependent. The up-front bonus payment $R(\theta)$ is used to capture the

remainder of the available surplus, net of information rents.

E Computational Model Details

The model has a number of state variables. The first is the firm's type, θ , known to the firm but not to the mineral owner. The second is the contract characteristics χ_i , as well as which stage i the lease is in. The third is the time-specific value of prices P_t and the expected drilling and production costs C_t , which evolve stochastically. We use the subscript triplet $\{i, t, \theta\}$ to tersely denote the state variables, e.g., $V_{i,t,\theta}^a$ denotes the expected value to the firm of taking action a given lease regime χ_i , time period t , current prices and costs P_t and C_t , and firm type θ .

The model is solved with backward induction. Therefore we discuss the solution in reverse order

E.0.1 Extension Phase ($t \geq \bar{T} + 1$, $i = 2$): The Firm's Problem

In the extension period, the firm can drill ($a = 1$) or not drill ($a = 0$). The expected payoff from drilling in period t and in lease regime $i = 2$, with firm type θ , is determined by contract terms k_i and S_i , the firm's expected present value of production θ , the natural gas price P_t , and the drilling and extraction cost C_t :

$$V_{i,t,\theta}^1 = (1 - k_i)P_t\theta_t - C_t + S_i \quad (22)$$

If the firm does not drill ($a = 0$), it continues in the game and receives the continuation value:

$$V_{i,t,\theta}^0 = \delta E[V_{i,t+1,\theta}] \quad (23)$$

The continuation value $E[V_{i,t,\theta}]$ is the expected value conditional on observable states and taken as an expectation over all possible draws of shocks. These shocks are denoted as $\nu_{t,\theta}^a$ for each action a , and they additively affect the payoff of each action. The firm's total payoff for choosing action a is $V_{i,t}^a + \nu_{t,\theta}^a$. Therefore the expected value for each time period is:

$$E[V_{i,t,\theta}] = E[\max\{V_{i,t,\theta}^1 + \nu_{t,\theta}^1, V_{i,t,\theta}^0 + \nu_{t,\theta}^0\}] \quad (24)$$

We assume that $\nu_{t,\theta}^1$ and $\nu_{t,\theta}^0$ are i.i.d. draws from a type 1 extreme value distribution with mean zero and scale parameter σ_ν . With these distributional assumptions on the shocks $\nu_{t,\theta}^a$, the probability of drilling at time period t conditional on not having drilled before and given lease state i and firm type θ is:

$$p_{i,t \neq \bar{T}, t > 0, \theta}^d = \frac{\exp\left(\frac{V_{i,t,\theta}^1}{\sigma_\nu}\right)}{\exp\left(\frac{V_{i,t,\theta}^1}{\sigma_\nu}\right) + \exp\left(\frac{V_{i,t,\theta}^0}{\sigma_\nu}\right)} \quad (25)$$

We solve for the value functions and drilling probabilities using value function iteration.

E.0.2 End of Primary Term ($t = \bar{T}$, $i = 1$): The Firm's Problem:

In $t = \bar{T}$, the firm's problem is very similar to the case of $t \geq \bar{T} + 1$, with two exceptions. The first is that the lease is in the last phase of the primary term where $i = 1$, meaning that if the firm drills, its payoffs depend on χ_1 rather than χ_2 . The second is that rather than choosing between drilling and waiting, the firm instead chooses between drilling, abandoning the lease, and signing the extension. We assume that the random shock for abandoning is the same as the random shock for extending and is equal to $\nu_{i,t,\theta}^0$. Therefore we can treat this as a nested problem where the firm first chooses whether to drill or not drill. Then conditional on not drilling, the firm must either decide whether to abandon or to extend. As above, the action to not drill is denoted as $a = 0$ and the value of not drilling is:

$$V_{i=1,t=\bar{T},\theta}^0 = \max\{-R_2 + \delta E[V_{i=2,\bar{T}+1,\theta}], 0\} \quad (26)$$

The firm's expected value is the maximum expected value over the choices to drill and not drill:

$$E[V_{i,t,\theta}|t = \bar{T}, i = 1] = E[\max\{V_{i,\bar{T},\theta}^1 + \nu_{i,\theta}^1, V_{i,\bar{T},\theta}^0 + \nu_{i,\theta}^0\}] \quad (27)$$

With these distributional assumptions on the shocks $\nu_{i,t,\theta}^a$, the probability of drilling in period $t = \bar{T}$, given that it has not drilled previously, and given lease state i and firm type θ is:

$$p_{i=1,t=\bar{T},\theta}^d = \frac{\exp\left(\frac{V_{i=1,\bar{T},\theta}^1}{\sigma_\nu}\right)}{\exp\left(\frac{V_{i=1,\bar{T},\theta}^1}{\sigma_\nu}\right) + \exp\left(\frac{V_{i=1,\bar{T},\theta}^0}{\sigma_\nu}\right)} \quad (28)$$

Similarly, the probability of extending, given that it has not drilled previously, can be written as below, where the indicator term indicates whether the firm would prefer to extend rather than abandon:

$$p_{i=1,t=\bar{T},\theta}^e = \frac{\exp\left(\frac{V_{i=1,\bar{T},\theta}^0}{\sigma_\nu}\right)}{\exp\left(\frac{V_{i=1,\bar{T},\theta}^1}{\sigma_\nu}\right) + \exp\left(\frac{V_{i=1,\bar{T},\theta}^0}{\sigma_\nu}\right)} \cdot 1\left[V_{i=1,\bar{T},\theta}^0 \geq 0\right] \quad (29)$$

E.0.3 End of Primary Term ($t = \bar{T}$, $i = 1$): The Mineral Owner's Problem:

When deciding which contract terms to offer for the extension, the mineral owner predicts how the firm will respond to its choice of contract terms χ_2 . For every possible price path, for each period $t \geq \bar{T} + 1$, and for each type θ , the mineral owner forecasts the probability of each of the firm's actions a and the effect of each action on the mineral owner's payoff. Specifically, if firm type θ drills ($a = 1$) in period t in lease regime $i = 2$, the mineral owner receives a payment in period t of $k_2 P_t \theta_t - S_2$. If the firm continues in the lease, the mineral

owner receives zero in that period. Using the drilling probabilities from equation (25) and the state transition matrix for P_t and C_t , the owner can calculate its expected present value payment for each firm type θ , given contract terms χ_2 .

As discussed in the main text, we assume that the royalty k_2 and subsidy S_2 are the same as those from the primary term, consistent with practice in the Haynesville. The owner's problem is then simply to choose the optimal extension bonus R_2 . Because the owner can calculate both its own expected value from the continuation period as well as the expected values $V_{i=1, t=\bar{T}, \theta}^0$ for each type θ , this problem is essentially an optimal reserve price problem. As noted in the main text, we make the tractability assumption that the distribution of types that the owner believes it faces at \bar{T} is the same as the initial distribution at $t = 0$.

E.0.4 Primary term $1 \leq t \leq \bar{T} - 1$: Firm's problem:

During the periods where $1 \leq t \leq \bar{T} - 1$, the firm can choose between drilling ($a = 1$) and waiting ($a = 0$). We can write the same value functions and choice probabilities as above in subsection E.0.1. The only difference is that here the lease regime is $i = 1$ rather than $i = 2$, meaning the payoffs from drilling depend on contract terms χ_1 . The payoffs of waiting and therefore overall payoffs also depend on the firm's expectations of how the mineral owner will choose χ_2 (and in particular R_2).

We solve for the firm's drilling probabilities and value function in each period of the primary term by iterating backwards from the values calculated in subsection E.0.2 and the optimal extension bonus R_2 calculated in subsection E.0.3.

E.0.5 Initial time period $t = 0$: Firm's problem

Given the initial contract terms χ_1 that the mineral owner offers the firm, the firm will choose to accept and pay the initial bonus R_1 if the expected profits exceed the bonus. Therefore the firm accepts if:

$$\delta E[V_{1,1,\theta}] \geq R_1 \tag{30}$$

We assume that there are no shocks here, which implies that there will be a threshold value $\theta^*(\chi_1)$ such that $\delta E[V_{1,1,\theta^*(\chi_1)}] = R_1$ and all firm types $\theta \geq \theta^*(\chi_1)$ will accept the lease.

E.0.6 Initial time period $t = 0$: Mineral Owner's problem

In the initial period, the mineral owner chooses contract term χ_1 , including primary term \bar{T} , to maximize the expected sum of payoffs during the primary term, from any extension payments, and payoffs during the extension.

F Additional Wells

In some analyses we augment the model so that the lease can accommodate multiple wells. Consistent with industry practice, we assume the firm needs to only drill one well to hold the lease. Once one well is drilled, the lease is held by production, and additional drilling

can happen at any later time period. We assume that the total number of wells that can be drilled on the lease is M .

To simplify the analysis, we make some assumptions on the functional form of payoffs. We assume that there is an additive shock $\nu_{t,\theta}^1$ that applies to each well that the firm drills in period t and another additive shock $\nu_{t,\theta}^0$ that applies to each (remaining) possible well that the firm could drill but does not drill in period t . We also assume no economies of scale in drilling costs and that any cost subsidy S_i is paid for each well. Therefore the total one-time payoff for drilling N wells out of a total of M possible wells is $N[(1 - k_i)P_t\theta_i - C_t + S_i + \nu_{t,\theta}^1] + (M - N)\nu_{t,\theta}^0$.

These additivity assumptions allow us to significantly decrease the number of choices we need to consider: they imply that if a firm prefers to drill a positive number of wells, it will prefer to drill all of the wells, because profits multiplicatively increase in the number of wells drilled. Therefore, in examining an infinite horizon game, we need only examine the choice to drill zero wells or to drill all remaining possible wells. In addition, if the lease is in its primary term and is not held by production, we must also consider the possibility that the firm would prefer to drill a single well: drilling a single well switches the lease regime from finite horizon to infinite horizon, which is valuable because the firm would have to pay an extension bonus if it instead wished to continue into the infinite horizon without drilling.

In this model, we use the tilde to denote the firm's per-well continuation value, equal to the total continuation value divided by the number of total remaining wells that can be drilled on the lease. We use $E[\tilde{V}_{i,t+1,\theta}]$ to denote the firm's per-value continuation value at t if the lease has not been held by production at period t . We use $E[\tilde{W}_{i,t+1,\theta}]$ to denote the firm's per-well continuation value if the lease has been held by production by date t (e.g., if there has been drilling at t or before). Because the constraint of a primary term decreases the expected continuation value of the lease, $E[\tilde{V}_{i,t+1,\theta}] < E[\tilde{W}_{i,t+1,\theta}]$.

We focus here on the case where $t \in [1, \bar{T} - 1]$ and for the period where the lease is in its primary term (where no drilling has been done). As discussed above, the firm has three choices. First, the firm may choose to drill all M wells ($a = M$), thus ending the optimal stopping problem. The firm receives M times the per-well profits.⁴¹

$$V_{i,t,\theta}^M = M[(1 - k_i)P_t\theta_i - C_t + S_i] \quad (31)$$

Second, the firm may choose to drill only one well ($a = 1$), thus holding the lease. As in the single-well case, the firm receives the profits from that decision. In addition, the firm receives the continuation payoff associated with the option to drill $M - 1$ future wells:

$$V_{i,t,\theta}^1 = (1 - k_i)P_t\theta_i - C_t + S_i + \delta(M - 1)E[\tilde{W}_{i,t+1,\theta}] \quad (32)$$

Finally, the firm can retain the lease without drilling ($a = 0$). It gets the continuation value associated with an undrilled lease capable of holding M wells.

$$V_{i,t,\theta}^0 = \delta ME[\tilde{V}_{i,t+1,\theta}] \quad (33)$$

And, as discussed above, we also assume additive shocks to the firm's payoff: for each well that it drills, it gets the shock $\nu_{t,\theta}^1$. For each potential well that it does not drill, it gets

⁴¹Again, we simplify the timing of drilling by assuming all extraction occurs immediately.

$\nu_{t,\theta}^0$. Thus the full payoffs from each of the three choices are:

$$\begin{aligned}
&\text{Drill all } M \text{ wells: } V_{i,t,\theta}^M + M\nu_{t,\theta}^1 \\
&\text{Drill 1 well: } V_{i,t,\theta}^1 + \nu_{t,\theta}^1 + (M-1)\nu_{t,\theta}^0 \\
&\text{Drill zero wells (continue): } V_{i,t,\theta}^0 + M\nu_{t,\theta}^0
\end{aligned} \tag{34}$$

This structure leads to a tractable set of choice probabilities. Let $u_{i,t,\theta}^1 = (1 - k_i)P_i\theta_i - C_t + S_i$ denote the deterministic payoff from drilling a single well. Combining equation (34) with equations (31)–(33), we can characterize the firm's choice probabilities. The firm prefers drilling all M wells to a single well if:

$$\nu_{t,\theta}^1 - \nu_{t,\theta}^0 > \delta E[\tilde{W}_{i,t+1,\theta}] - u_{i,t,\theta}^1 \tag{35}$$

In addition, the firm prefers drilling M wells to not drilling any wells if:

$$\nu_{t,\theta}^1 - \nu_{t,\theta}^0 > \delta E[\tilde{V}_{i,t+1,\theta}] - u_{i,t,\theta}^1 \tag{36}$$

Since $E[\tilde{W}_{i,t+1,\theta}] > E[\tilde{V}_{i,t+1,\theta}]$, it is clear that if the firm prefers drilling M wells to drilling one well, it also prefers drilling M wells to drilling zero wells.

Finally, the firm prefers drilling one well to drilling no wells if:

$$\nu_{t,\theta}^1 - \nu_{t,\theta}^0 > M\delta E[\tilde{V}_{i,t+1,\theta}] - (M-1)\delta E[\tilde{W}_{i,t+1,\theta}] - u_{i,t,\theta}^1 \tag{37}$$

This system of preferences is then an ordered logit. We have the following choice probabilities:

$$\begin{aligned}
p_{i,t,\theta}^M &= 1 - \frac{1}{1 + \exp\left\{\frac{1}{\sigma_\nu} \left(u_{i,t,\theta}^1 - \delta E[\tilde{W}_{i,t+1,\theta}]\right)\right\}} \\
p_{i,t,\theta}^1 &= \frac{1}{1 + \exp\left\{\frac{1}{\sigma_\nu} \left(u_{i,t,\theta}^1 - \delta E[\tilde{W}_{i,t+1,\theta}]\right)\right\}} \\
&\quad - \frac{1}{1 + \exp\left\{\frac{1}{\sigma_\nu} \left((M-1)\delta E[\tilde{W}_{i,t+1,\theta}] + u_{i,t,\theta}^1 - M\delta E[\tilde{V}_{i,t+1,\theta}]\right)\right\}} \\
p_{i,t,\theta}^0 &= \frac{1}{1 + \exp\left\{\frac{1}{\sigma_\nu} \left((M-1)\delta E[\tilde{W}_{i,t+1,\theta}] + u_{i,t,\theta}^1 - M\delta E[\tilde{V}_{i,t+1,\theta}]\right)\right\}}
\end{aligned} \tag{38}$$

There are three distinct intervals on the support of $\nu_{t,\theta}^1 - \nu_{t,\theta}^0$, each associated with a single optimal action. This result produces some intuitive firm behavior. Suppose that the lease is early in the primary term. Then $E[\tilde{W}_{i,t+1,\theta}]$ and $E[\tilde{V}_{i,t+1,\theta}]$ are reasonably close to one another because the value of moving into the indefinite secondary term immediately is low (if they were indeed equal, the interval on which $a = 1$ is chosen would be zero). Thus, conditional on wanting to drill at least one well early on in the lease, the firm is likely to drill

all of them. On the other hand, suppose $E[\tilde{W}_{i,t+1,\theta}]$ is much larger than $E[\tilde{V}_{i,t+1,\theta}]$ because the lease is late in the primary term. Then the payoff from moving into the secondary term is large, and the firm will be more likely to drill just one well conditional on drilling any at all.

The math is somewhat more complicated when we move to the case where the primary term is about to expire ($t = \bar{T}$). Now the firm's choice to drill zero wells now includes two possible sub-actions: pay to extend, or abandon. Now we write the value of drilling zero wells as:

$$V_{i,t=\bar{T},\theta}^0 = \max\{-R_2 + \delta ME[\tilde{V}_{i,t+1,\theta}], 0\} \quad (39)$$

Again, this can be solved as an ordered logit, where the firm chooses between drilling all M wells, drilling one well, or drilling zero wells. Then, conditional on drilling zero wells, the firm continues only if $V_{i,t=\bar{T},\theta}^0 \geq 0$.

G Well decline estimation

To estimate well production decline, we follow Patzek et al. (2013), who derive decline curves for shale gas formations. They show that initial production declines slowly and inversely proportional to the square root of time, but then begins to decline at an exponential rate.⁴² To capture this, we assume that cumulative production of natural gas for well i at month t takes the following functional form:

$$m_i(t) = \begin{cases} M_i \sqrt{t/\tau} & \text{if } 0 \leq t \leq \tau \\ M_i + \frac{M_i}{2\tau d} [1 - \exp(-d(t - \tau))] & \text{if } t > \tau \end{cases} \quad (40)$$

where τ is the time at which the decline function changes due to cross-fracture interference, d is the exponential decline rate of production after interference begins, and M_i is a well-specific production multiplier corresponding to the expected cumulative production at $t = \tau$.⁴³ Production is measured in thousands of cubic feet (Mcf).

Before estimating these parameters, we make a number of adjustments to the data. First, because horizontal production is substantially affected by the length of the lateral well leg, we normalize the measure of cumulative production by a scalar s_i which is equal to 1485 meters divided by the length of the lateral portion of the well. We drop any well with missing lateral length information or a well lateral of less than 150 meters to eliminate potentially misclassified vertical wells.⁴⁴

We find that about 3% of wells had recompletions. As recompletions are designed to rapidly increase production, we exclude from the data observations that come during or after months where well recompletions were performed. Later on, when we simulate the expected productivity of wells, we assume no recompletions.

⁴²We have also tried estimating the initial rate of decline, and get similar results.

⁴³This functional form incorporates the structure of Patzek et al. (2013), but also imposes the requirement that cumulative production be differentiable at $t = \tau$.

⁴⁴This is a reasonable breakpoint: there are no wells with laterals between 150 and 400 meters in our data.

Following Patzek et al. (2013), we limit the sample to observations that are the fourth month or later ($t \geq 4$) because early months of production tend to be very noisy. This is due in part to the fact that hydrofracturing water is still being back-produced in early months of production. Similarly, for any given date with no production but where there is production both before and after the date, we assume that the production process is paused on that date and resumes when production resumes.

Rather than estimating τ directly, we use estimates from Male et al. (2015) for the Haynesville, who find that τ is 14.16 months. We then use non-linear least squares to find the values of M_i and d that minimize the sum of the squared differences between true and predicted log cumulative production, as shown in equation 41.

$$\sum_i \sum_{t|t \geq 4} (\log m_i(t) s_i - \log \widehat{m}_i(t|M_i, \tau, d))^2 \quad (41)$$

The estimated decline parameter d is equal to 0.0361. The 25th, 50th, and 75th percentiles of the estimated M_i are 1.47 million, 2.01 million, and 2.56 million, respectively.

We then use our estimates of d and M_i to predict total discounted well production (Equation 42). Following Gülen et al. (2015), we assume that wells have a total production lifetime of 20 years. We use a monthly discount factor of 0.9971, which corresponds to an annual discount factor of 0.909. Figure 3, Panel (A) shows a map of our measures of present value of total well production. Where there are multiple wells, we take an average over all wells within the unit. Sections with no drilling have no shading and are labeled NA.

$$Y_i = \sum_{t=1}^{240} [\widehat{m}_i(t|M_i, \tau, d) - \widehat{m}_i(t-1|M_i, \tau, d)] \delta^{t-1} \quad (42)$$

H Geographic smoothing

Our next step is to compute measures of expected productivity which we do through geographic smoothing of using nearby values of production to impute expected production. We do this for two reasons: The first is that some units did not have drilling, and so this method allows us to predict expected production. The second is that our measures of productivity Y_i are noisy, and this smoothing gives a more reasonable measure of a firm's expected productivity.

We do this in three steps. In the first step, we find an optimal spatial bandwidth that best predicts well productivity using neighboring unit productivity. Then in the second step, we estimate a semiparametric model that allows for a flexible time trend in productivity (Robinson 1988). Finally, in the third step, we net out the effects of the time trend and impute the expected production of a typical well within any given unit.

For our first step, we use the logarithm of our total production predictions $\ln(Y_i)$. We use a Gaussian kernel of functional form $\phi(0, \sigma_\phi^2)$, where ϕ is the pdf of the normal distribution with a mean of 0 and a variance of σ_ϕ^2 . We use a weighted average of all wells i' that are not in i 's unit to predict production of well i . The optimal bandwidth is the value that minimizes the sum of the squares of the differences between actual production and predicted production:

$$\sigma_\phi^2 = \arg \min_{\sigma^2} \sum_i \left(\ln(Y_i) - \frac{\sum_{i' \notin u(i)} \phi(d(i, i'), \sigma^2) \cdot \ln(Y_{i'})}{\sum_{i' \notin u(i)} \phi(d(i, i'), \sigma^2)} \right)^2 \quad (43)$$

Here, i and i' index wells, $d(i, i')$ is the distance between wells i and i' , and $\ln(Y_i)$ is the log production of well i .

Then in our second step, we implement the difference estimator suggested in Robinson (1988).⁴⁵ We model productivity as:

$$\ln(Y_i) = g(lon_i, lat_i) + \sum_s \beta_s X_s(t_i) + u_i \quad (44)$$

where the $X_s(t_i)$'s are a vector of restricted cubic spline basis functions of the well completion date t_i and $g(lon_i, lat_i)$ is a nonparametric function of well coordinates $\{lon_i, lat_i\}$.⁴⁶

First, we use our estimated optimal bandwidth σ_ϕ^2 to predict productivity at the site of each well (m_i^y) and compute residuals (e_i^y). Second, we use the same approach to geographically smooth each of the basis functions that make up the completion date spline and obtain predictions ($\mathbf{m}_i^{\mathbf{X}}$) and residuals ($\mathbf{e}_i^{\mathbf{X}}$). Third, we regress e_i^y on $\mathbf{e}_i^{\mathbf{X}}$. This returns an estimate of the spline coefficients $\hat{\beta}_s$ that account for the fact that different areas were drilled at different times. This spline function is shown in Figure [[XX]]. Fourth, we subtract the product of the smoothed time basis functions and the spline coefficients from each well's smoothed predicted productivity:

$$\hat{m}_i = m_i^y - \sum_s \hat{\beta}_s m_i^{X_s} \quad (45)$$

We nonparametrically smooth these values to predict the productivity of a well drilled at the centroid of unit u , again using a bandwidth of σ_ϕ^2 :

$$\hat{g}(lat_u, lon_u) = \frac{\sum_{i' \notin u} \phi(d(u, i'), \sigma_\phi^2) \cdot \hat{m}_i}{\sum_{i' \notin u} \phi(d(u, i'), \sigma_\phi^2)} \quad (46)$$

Here u indexes units and $d(i', u)$ is a measure of distance between a well i and the centroid of a unit u . Finally, we shift the predicted productivity by the value of the time trend at the earliest completion date in our sample (Oct. 10, 2007) to model productivity at that point in time.

$$E[\ln(Y_u)|t_u = \text{Oct. 10, 2017}] = \hat{g}(lat_u, lon_u) + \sum_s \hat{\beta}_s X_s(\text{Oct. 10, 2007}) \quad (47)$$

[[Spline function plot to go about here.]]

⁴⁵In each nonparametric smoothing step, we use a Gaussian kernel with the same bandwidth σ_ϕ^2 .

⁴⁶We use 5 knots located at evenly-spaced quantiles of the completion date distribution.