Drilling Deadlines and Oil and Gas Development

Evan Herrnstadt
Ryan Kellogg
Eric Lewis*

October 5, 2023

Abstract

Oil and gas leases between mineral owners and extraction firms typically specify a date by which the firm must either drill a well or lose the lease. These deadlines are known as primary terms. Using data from the Louisiana shale boom, we first show that well drilling is substantially bunched just before the primary term deadline. This bunching is not necessarily surplus-reducing: using an estimated model of firms' drilling and input choices, we show that primary terms can increase total surplus by countering the effects of leases' royalties, as royalties are a tax on revenue and delay drilling. These benefits are reduced, however, when production outcomes are sensitive to drilling inputs and when drilling one well indefinitely extends the period of time during which additional wells may be drilled. We enrich the model to consider mineral owners' lease offers and find small effects of primary terms on owners' revenue.

*Herrnstadt: Congressional Budget Office; evan.herrnstadt@cbo.gov. Kellogg: Harris Public Policy, University of Chicago; and National Bureau of Economic Research; kelloggr@uchicago.edu. Lewis: Bush School of Government and Public Service, Texas A&M University; ericlewis@tamu.edu. We thank William Patterson, Pengyu Ren, Grant Strickler, and especially Nadia Lucas for outstanding research assistance. We are grateful for the many thoughtful suggestions from various conference and seminar participants. The analysis and conclusions expressed herein are solely those of the authors and do not represent the views of the U.S. Congressional Budget Office.

1 Introduction

Owners of subterranean oil and gas typically write contracts with specialized extraction firms to act as their agents because they lack the relevant expertise or capital necessary to extract their resources. In the United States, as well as several other countries, these contracts take the form of mineral leases that ubiquitously contain deadlines known as "primary terms". A primary term is a period of time during which the firm must drill at least one producing well. Drilling effectively extends the lease term until production ends; not drilling within the primary term ends the lease. Leases also typically specify a royalty payment from the firm to the owner. In this paper, we aim to understand the effects of this contract structure on drilling, production, and surplus outcomes. We emphasize the economic impacts of primary terms and how these effects are moderated by the distinctive characteristics of modern shale resources.

An oil and gas lease grants a firm an option, but not an obligation, to develop the mineral owner's property by drilling wells and extracting the hydrocarbons. Upon signing a lease, the firm pays the owner a flat fee, known as a "bonus". The primary term specifies a period of time (typically 3 to 10 years) that the firm has to drill at least one well and commence production. If it does so, the lease is then "held by production" and remains in effect until the firm ceases production. The firm may then also drill additional wells on the parcel to increase its overall production rate. On the other hand, if the firm does not complete a well by the end of the primary term, the lease terminates, and the mineral owner is free to sign a new contract with another firm or re-contract with the original firm.

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. Royalties are often significant, as the royalty rate typically lies between 12.5% and 25%. Brown et al. (2016) estimates that royalty payments associated with the six largest U.S. shale plays totaled $39 billion in 2014.

The royalty and primary term clauses distort firms' incentives regarding when to drill wells and how much effort to invest in fracking and well completion. The incentive to drill at least one well before primary term expiration has received considerable attention within the industry, with numerous reports of firms drilling unprofitable wells for the sake of holding their lease acreage. For instance, the *San Antonio Express News* reported in 2012 that "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" (Hiller 2012). Although royalties are less prominent in the news, they also distort firms' decisions because they are a tax on revenue only, thereby driving a wedge between firms' profit and total surplus.

We begin our analysis by studying data from the Haynesville Shale in Louisiana, where the institutional structure and data availability are conducive for studying lease terms. We discuss relevant institutional features of the Haynesville in Section 2, discuss our data sources in Section 3, and then show in Section 4 that there is substantial bunching of drilling in the months just prior to lease expiration. We further show that many leases are characterized by having only a single well that was drilled just before lease expiration, suggesting that drilling in these areas was primarily motivated by holding acreage for future wells rather than by immediate profits.

While the bunching analysis suggests that the primary term influences firms' drilling decisions, it does not identify what drilling decisions, production, and surplus would be absent the primary term. The evident distortion might suggest that the primary term reduces surplus; however, primary terms act in the presence of a royalty rate that is typically 20-25%. This tax on revenue reduces firms' returns to drilling so that they require more favorable price and cost conditions to drill a well, delaying drilling in expectation relative to what would have occurred absent the royalty. The primary term may then increase total surplus by counteracting the royalty's drilling timing distortion.

To quantify these effects, in Section 5 we present an econometric model of the major decisions the firm makes after acquiring a lease. First, our model incorporates the decision of whether and when to drill. We model this decision as an optimal stopping problem in the spirit of Kellogg (2014), Bhattacharya et al. (2018), Ordin (2019), and Agerton (2020), wherein the firm chooses when to drill during the primary term in the presence of stochastic natural gas prices and drilling costs. Second, our model accounts for the firm's input quantity decision conditional on drilling, which is a new feature in this class of models. As discussed in Covert (2015), modern shale wells (unlike conventional wells) require large inputs of fracking fluid in order to produce commercial quantities of hydrocarbons. Exploiting variation in wells' use of water over time—which we take as driven by variation in gas prices and the prices of inputs into well drilling and completion—we estimate that the marginal productivity of water is large: moving from the 25th to the 75th percentile of water increases production by about one-half its interquartile range in the data. Third, in our model, like that in Agerton

(2020), the value of drilling the first well includes the value of unlocking an indefinite option to drill additional wells later.

We use our model to jointly estimate drilling costs and natural gas productivity across the Haynesville, allowing for both observable and unobservable (to the econometrician) productivity shifters that affect firms' drilling decisions and drilling outcomes. Our estimates imply that some Haynesville wells—especially those drilled just before significant lease acreage expired—had negative expected profits, consistent with the notion that they were drilled to preserve lease acreage and future option value.

In Section 6, we discuss counterfactuals in which Haynesville leases omit a primary term, royalty, or both. A lease with neither provision leads the firm to make drilling and input decisions that maximize total surplus. A primary term alone substantially accelerates drilling—with drilling probabilities peaking just prior to expiration—reducing surplus. But in the presence of a 25% royalty, which causes drilling to be delayed, we find that the acceleration of drilling under the primary term is surplus-increasing, on average. The surplus gain is modest, amounting to 7.3% of the surplus loss imposed by the 25% royalty. This limited efficacy arises from two factors that are important to the shale oil and gas setting. First, primary terms do not directly affect firms' water input choices, which are substantially distorted by the royalty. Second, primary terms hasten the timing of only the first well drilled on a lease, not any later wells. When we simulate a case that is more akin to non-shale, conventional oil and gas development—in which a lease can only accommodate one well and the marginal productivity of water is zero—we find that primary terms are more effective and recover 42.6% of the surplus lost by the 25% royalty.

Finally, in Section 7 we study how primary terms interact with royalties to affect the mineral owner's expected discounted revenue from a lease. We adopt a modeling framework in which firms have a hidden signal about productivity and owners can make take-it-or-leave-it contract offers, following the literature on oil and gas auctions (see Haile et al. (2010) and Kellogg and Reguant (2021) for reviews) and especially recent papers that study owner-optimal royalties in auctions for state-owned parcels (Bhattacharya et al. 2018; Ordin 2019; Kong et al. 2022). In these papers, a higher royalty rate trades off reductions in firms' information rents with decreases in firms' likelihood of drilling conditional on being awarded a lease, per theoretical arguments from Hendricks et al. (1993) and Skrzypacz (2013).

Our analysis in Section 7 builds on this previous work in two ways. First, we use our estimates of the marginal productivity of fracking inputs to highlight that, in the new era of shale oil and gas, the input choice distortion induced by the royalty can substantially reduce the owner's revenue-maximizing royalty rate. We find that the owner-optimal royalty is 25% in our baseline model but 39% in an alternative specification that sets the marginal productivity of water to zero. Second, we evaluate the impact of primary terms conditional on the royalty. We find that a primary term slightly decreases owner revenue in our baseline model (relative to a royalty-only contract) but slightly increases it when we shut down water inputs and allow only one well to be drilled on the lease. These findings are consistent with results from an analytically tractable version of our model, where we show that if the sensitivity of production to input choice is high, then provisions that induce the firm to accelerate drilling decrease the owner's expected revenue. Because the primary term more negatively affects the firm's value when productivity is low than when it is high, it also effectively increases dispersion across firm types, thus reducing the owner's ability to capture value. Moreover, the owner responds to this increased dispersion by setting the bonus so that fewer firm types agree to the lease offer, reducing total surplus.

We focus our attention on primary terms in the shale oil and gas industry because primary terms play an important and under-studied role in the development of shale resources—which now account for the majority of U.S. oil and gas production—and because this industry is rich in data. But the underlying economics are likely to be relevant in other settings in which resource owners sell time-limited development options to other agents. For instance, master franchise contracts in retail settings typically specify royalty payments to the franchisor and impose a finite time for the franchisee to develop a minimum number of units (Kalnins 2005). Licenses for adaptations of creative works (such as adaptations of novels for screenplays) often allow producers only a finite period to commence or complete production, lest the property rights revert to the original author (Litwak 2012). And U.S. Federal Communication Commission spectrum auctions impose buildout requirements upon winning firms (GAO 2014). Our hope is that this paper can serve as a springboard for studying the economics of contract term length in these and other settings.

2 Institutional background

The development of new techniques combining horizontal drilling with hydraulic fracturing in the early 2000s led to drilling booms in shale formations throughout the United States. We focus on the portion of the Haynesville Shale located in Louisiana for two reasons. First, the Haynesville produces almost exclusively dry natural gas, allowing us to focus our analysis on a single output. Second, the economic and legal institutions in Louisiana that shape 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. This section summarizes these institutions.

When a firm is interested in drilling on privately-owned land, it must negotiate a lease with the mineral owner.[1] U.S. oil and gas leases almost always include a cash bonus paid at signing, a royalty, and a primary term. The royalty rate specifies the fraction of oil and gas revenue that must be paid to the mineral owner, and the primary term sets the amount of time that the firm has an option to drill and commence production before it loses the lease. Once a productive well is drilled, the lease is "held by production", which means that the lease continues in force as long as the firm maintains commercial oil and gas production on the lease. A lease may also include an extension clause, which gives the firm an option to extend the lease for a pre-specified amount of time in exchange for an additional, pre-specified payment to the mineral owner.

In practice, leases typically have a clause that allows the firm to hold a lease beyond expiration even if it is not producing, so long as it is actively in the process of drilling or fracking a well. Our analysis will therefore focus on "spudding" a well—i.e., commencing drilling—as the necessary step to hold a lease.

Leases are often small relative to the area drained by a modern shale well, which may have a horizontal length of 5,000 feet or more. Therefore, state regulators have established rules for combining leases into pooling units. In Louisiana, the default pooling unit for the Haynesville Shale is the square-mile section from the Public Land Survey System (PLSS). Typically, multiple firms will hold leases within a given pooling unit, and drilling operations then 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 lease holding firms on an acreage-weighted basis. Each firm then

[1] The lease structure we describe here also applies to publicly-owned oil and gas in the U.S.

distributes royalties on revenues to its mineral owners on an acreage-weighted basis.

Drilling a Haynesville well within a Haynesville pooling unit holds all current leases within the unit, not just those overlying the well itself. In addition, because horizontal wells in shale formations primarily recover gas that is located in rock close to the well bore, square-mile units have space for multiple horizontal wells that run parallel to one another. Thus, drilling a single well in a unit grants the operating firm the indefinite right to drill additional wells within the same unit.

Owners of minerals 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.[2] It is therefore the threat that acreage in a unit will convert from leased to unleased that gives firms an incentive to drill prior to the expiration of primary terms. A unit typically consists of many leases, not all of which expire at the same time. The drilling incentive provided by a given lease's pending expiration depends on the acreage of that particular lease as well as the schedule of expiration dates for remaining leases.

3 Data sample and summary statistics

This section summarizes our data on natural gas prices, rig dayrates, wells, leases, and units. We provide additional detail in Appendix A.

3.1 Price and rig dayrate data

Our measure of the price of natural gas is the 12-month natural gas futures price for delivery at Henry Hub, Louisiana, obtained from Bloomberg (2017).[3] For the period from 2009–2013 (during which most of the Haynesville drilling happens), the average natural gas price is \$5.07 per mmBtu (mmBtu = million British thermal units), with a minimum of \$3.39 and a maximum of \$7.75.[4] We also obtained data on rig dayrates, which are the cost of renting

[2]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"). Consequently, firms cannot earn strictly positive profits from unleased acreage.

[3]We use prices for delivery at a 12-month horizon because wells produce gas gradually rather than instantaneously, and 12 months is the longest horizon at which futures are consistently liquidly traded.

[4]We deflate all gas price, rig dayrate, and drilling cost data to December 2014 dollars using the Bureau of Labor Statistics' Consumer Price Index for all goods less energy, all urban consumers, and not seasonally

Table I: Summary statistics for wells

Variable	Obs	Mean	Std. Dev.	P10	P50	P90
Well spud year	2685	2010.5	1.5	2009	2010	2013
Well completion year	2685	2011	1.6	2009	2011	2013
Accounting well cost (millions, Dec 2014$)	2495	10.4	2.4	7.8	10.1	13.3
Water volume (millions of gallons)	2401	6	2.8	3.5	5.5	8.9
PV total production (millions mmBtu)	2484	3.6	1.5	1.8	3.5	5.4

Note: The descriptive statistics in this table include all Haynesville wells, as defined in Section 3.2 and Appendix A.2. P10, P50, and P90 refer to the 10th, 50th, and 90th percentile of the relevant variable. The number of observations varies across rows because some variables are missing for some wells.

a drilling rig for one day, from Enverus (2017). The average dayrate from 2009–2013 was $16,841, with a minimum of $12,470 and a maximum of $18,721.

3.2 Well data

We obtained data on well drilling and completions from Enverus (2016a), Louisiana Department of Natural Resources (2016a), Louisiana Department of Natural Resources (2016b), and Louisiana Department of Natural Resources (2016c). These data include permit dates, spud dates, completion dates, the volume of water used in hydraulic fracturing, whether the well targets the Haynesville formation, and drilling and completion costs reported to the Louisiana DNR. We obtain well-level monthly production data from Enverus (2016b).

We focus our analysis on wells that targeted the Haynesville formation. In Table I, we present summary statistics for these wells. Most Haynesville wells were spudded and completed between 2009 and 2013. Water used in hydraulic fracturing ranged from less than 3.5 million gallons to more than 8.9 million gallons. Reported drilling and completion costs range from less than $7.8 million to more than $13.3 million.

To estimate the cumulative lifetime production from each well, we fit a decline curve to Enverus's monthly well-level production data.[5] Our decline model, which we discuss in Appendix A.2, is based on the functional form derived in Patzek et al. (2013). We use the estimated parameters to predict well-level production over time (extrapolating beyond our observed data) and then to predict the present value of each well's total lifetime cumulative

adjusted (Bureau of Labor Statistics 2018). The CPI series ID is CUUR0000SA0LE.

[5]Following Anderson et al. (2018), and consistent with Newell et al. (2019), we assume that wells' production decline rate is unaffected by natural gas price shocks.

Figure 1: Haynesville unit-level productivity

(a) Actual average log production per well

(b) Original gas in place (OGIP)

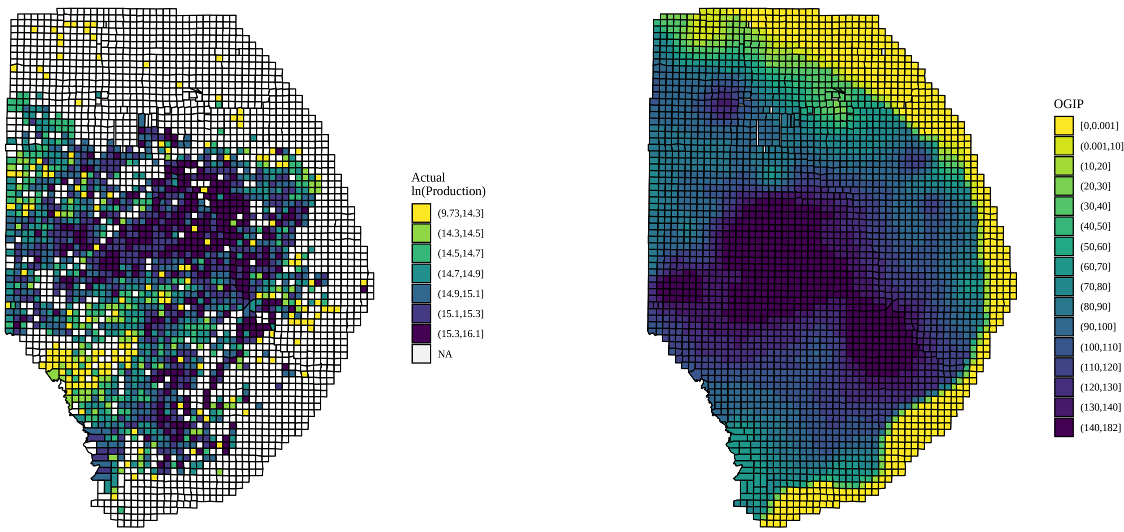

Note: Panel (a) is a map of Haynesville units showing the log of the calculated present value of aggregate well production, averaged within each unit, using decline estimation procedures discussed in Section 3 and Appendix A.2. In Panel (b), we plot original gas in place (OGIP), from Gülen et al. (2015).

production. Production summary statistics are shown in the last row of Table I. We find that the median present value of cumulative production is 3.5 million mmBtu, with 10th and 90th percentiles of 1.8 million mmBtu and 5.4 million mmBtu, respectively. In Figure 1, we show the average log of estimated productivity for all wells in our sample.

3.3 Lease data

We compile data from Enverus (2016a) on the universe of oil and gas leases in Louisiana that started between 2002 and 2015. These data include the start date of the lease, the primary term, any extension options, the royalty rate, the lease's PLSS section, and the acreage of the lease. The initial signing bonus is not recorded because state and local recorders do not require it, and because firms typically wish to keep lease terms confidential.

We focus on leases that are within our sample of Haynesville pooling units, as described in Section 3.4 below. In Table II, we present descriptive statistics for the 38,694 leases in this sample. Leases typically started between 2005 and 2011. Leases range from less than 0.20 acres to more than 240 acres, with a mean of about 80 acres. Typical royalty rates

Table II: Summary statistics for leases

Variable	Obs	Mean	Std. Dev.	P5	P50	P95
Year lease starts	38694	2008.5	1.7	2005	2008	2011
Year lease ends	38694	2011.5	1.8	2008	2011	2014
Primary term length (months)	38694	37.1	6.3	36	36	60
Indicator: Has extension clause	38612	.8	.4	0	1	1
Extension length (months)	29973	24.1	2.8	24	24	24
Royalty rate	30215	23	2.9	18.8	25	25
Area in acres	38428	79.9	386.7	.2	5.4	244.1

Note: The descriptive statistics in this table include all leases associated with Haynesville units included in our analysis sample, as defined in Section 3.4. P5, P50, and P95 refer to the 5th, 50th, and 95th percentile of the relevant variable. The number of observations varies across rows because variables may be missing for some leases. The statistics for "extension length" are computed for the subsample of leases that have an extension clause.

are 25% (65% of leases), 20% (19% of leases), and 18.75% (11% of leases). About 92% of leases have 36-month primary terms; a small fraction of leases have 60 month primary terms. About 78% of leases have extension clauses, with the vast majority of extensions lasting 2 years.[6] Exercising the extension option requires the payment of an additional bonus, but these bonuses are not usually recorded in the lease documents.

3.4 Pooling unit data

We obtain shapefiles for designated Haynesville units from Louisiana Department of Natural Resources (2016a). These units are typically PLSS square-mile (640 acre) sections, though some units have slight irregularities. In addition to these DNR-designated Haynesville units, we also include in our sample PLSS sections that that lie within the convex hull of the DNR-designated units.

Since we are interested in how the incentive to hold acreage affects the drilling of Haynesville wells, we remove from our sample units that may be held by drilling or production from other oil and gas formations. We do so by dropping units that have leases executed prior to 2004, non-zero oil or gas production in 2006, or non-Haynesville wells drilled after 2000. Our remaining sample, which we refer to as our *analysis sample* because we use it for

[6] We show in Table A.II in Appendix B.2 that the correlations of royalties and primary terms with observable geologic quality and with natural gas prices are economically small. Extension clauses are associated with lower geologic quality units and lower gas prices.

Figure 2: Map of Louisiana Haynesville units

(a) Units in analysis sample **(b)** Wells drilled per unit

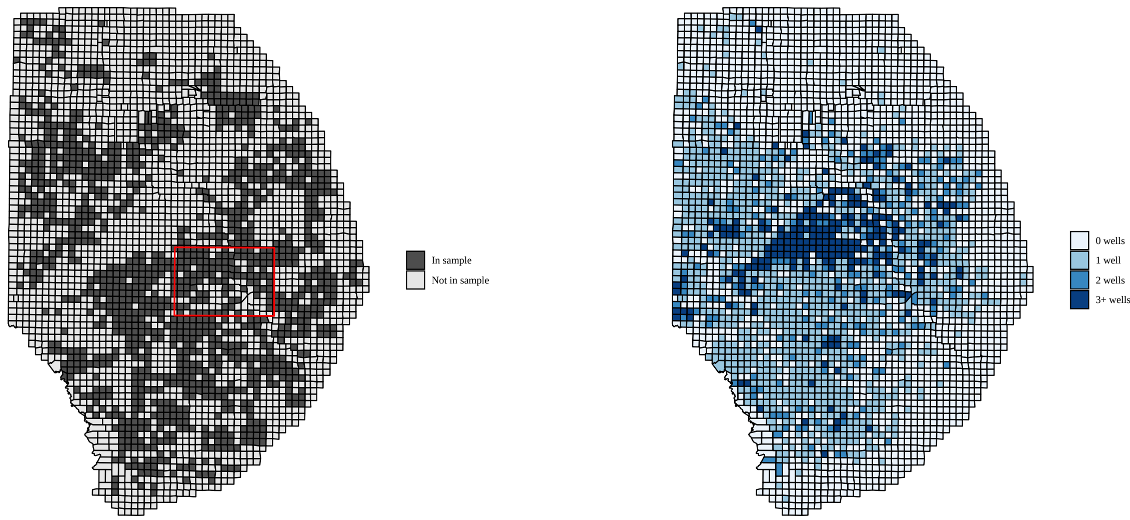

Note: Panel (a) is a map of Haynesville units (each square is a unit), where units that are in the analysis sample are colored dark. The rectangle is the outline of the map in figure 3. Panel (b) is a map of Haynesville units, with units colored by how many Haynesville wells were drilled as of March 2017.

the analysis in Section 4 below, includes 1,226 units, which we map in Panel (a) of Figure 2.

We match leases to units using the reported section in each lease document. In some cases, the reported total acreage of all leases in a unit exceeds the actual unit acreage due to likely duplicates in the data. We use a clustering procedure—which we describe in detail in Appendix A.3—to identify and downscale acreage for these likely duplicates.

To match wells to units, we use GIS techniques to identify which unit the majority of a well's horizontal leg passes through. Further details are in Appendix A.4. In Figure 3, we present a March 2017 snapshot of well laterals and pooling units for a selected portion of the Haynesville, illustrating the mapping of wells to units.

In Table III, we show summary statistics for our analysis sample of units. Units tend to have their first lease expire between 2008 and 2011, with a median of 2009. A total of 712 units (58%) have Haynesville wells drilled, with the first spud typically between 2008 and 2011. Of the units with drilling, 74% have only one well drilled, 15% have 2 wells drilled, 5% have 3 wells drilled, and 7% have 4 or more wells drilled. The most wells we observe in a single unit is 18. In Panel (b) of Figure 2, we map the number of wells drilled per unit.

Figure 3: An example of drilling patterns in the Haynesville Shale.

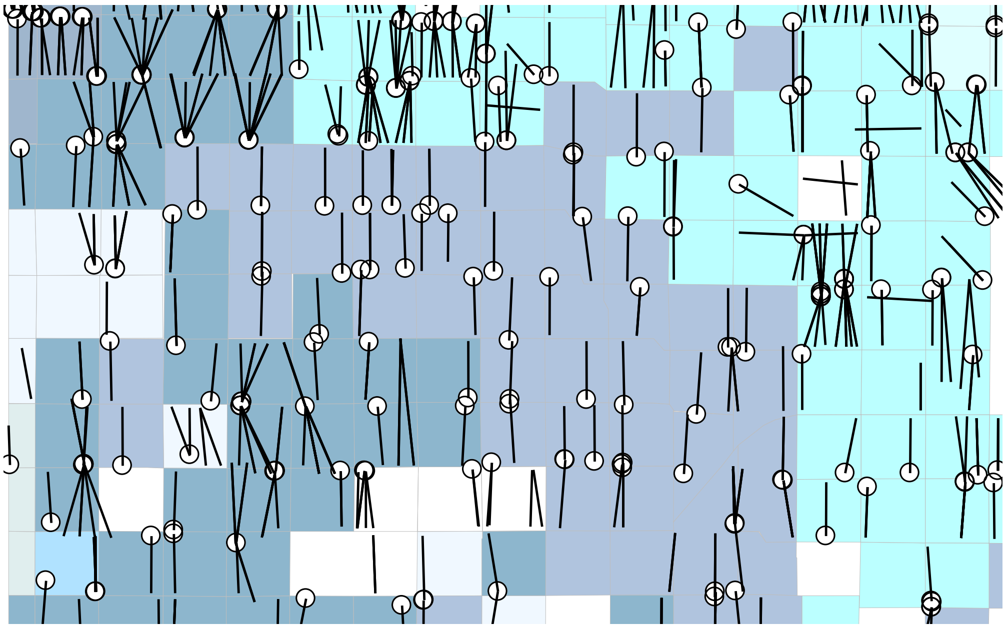

Note: Map produced using data from the Louisiana DNR's SONRIS. Each square is a unit, white dots are wellheads, and black lines are the approximate horizontal well path. Units are colored by unit operator. Data are as of March 2017 and include all Haynesville units, not just those in the analysis sample. The rectangle in Panel (a) of Figure 2 indicates the location of this example within the Haynesville Shale.

We also show summary statistics on geological "original gas in place" (OGIP) from Gülen et al. (2015) (also used in Agerton (2020)) in Table III. OGIP approximates total natural gas in the Haynesville as a function of formation thickness, porosity, temperature, and pressure, but not using information from production outcomes. We present a map of OGIP in Panel (b) of Figure 1. Taken together, the two panels depict the positive correlation between unit-level OGIP and production.

Finally, in Figure 4, we present time series aggregates, within our analysis sample of Haynesville units, for three major variables: the natural gas price, the number of leases signed, and the number of wells drilled. This figure shows that the gas price and Haynesville leasing peaked in early 2008, but drilling did not peak until about two years later, shortly before many leases were to expire. This pattern suggests that primary terms may have had a significant effect on aggregate drilling activity in the Haynesville, a possibility we examine more directly in our bunching analysis below.

Table III: Unit-level summary statistics

Variable	Obs	Mean	Std. Dev.	P5	P50	P95
Section acres	1226	641.7	13.7	620.2	642.7	662.3
Year first lease starts	1226	2006.5	1.4	2005	2006	2008
Year first lease expires	1226	2009.5	1.5	2008	2009	2011
Number of Hay. wells	712	1.6	1.6	1	1	4
Year of first Hay. spud	712	2009.8	1	2008	2010	2011
OGIP	1226	100.9	39.3	4.1	103.8	156.7

Note: The descriptive statistics in this table include all Haynesville units included in our analysis sample, as defined in Section 3.4. P5, P50, and P95 refer to the 5th, 50th, and 95th percentile of the relevant variable. The summary statistics in rows 4 and 5 are conditional on at least one well being drilled on the unit. OGIP denotes "original gas in place", as computed in Gülen et al. (2015), and is measured in billions of cubic feet of natural gas per square mile.

4 Evidence on primary terms and bunching of drilling

To study the role of lease expiration in motivating drilling in the Haynesville, we compare the date that the first Haynesville well is spudded in each unit in our analysis sample to the first date that a lease within the unit reaches the end of its primary term. In Panel (a) of Figure 5, we present a kernel-smoothed distribution of spud timing relative to that expiration date, along with a 95% confidence interval; in Panel (b) we present a histogram of the same data. The substantial spike in the density prior to the expiration of the first lease suggests that lease expiration is often a binding constraint. In Appendix B.1, we use a formal bunching test to confirm that this spike is large and statistically significant. Further, in Appendix B.2, Figure A.3, we show that wells drilled just before expiration are fully-completed, producing wells: they do not exhibit remarkably low production, water inputs, or reported drilling costs.

Some leases in the Haynesville have a built-in extension clause that allows the firm to pay an additional bonus to extend the primary term by two years. Accordingly, a secondary spike in drilling two years after the primary term expires is in evident both panels of Figure 5. In Panel (a) of Figure 6, we split our sample into units in which the first lease to expire had an extension clause versus units in which the first lease did not have such a clause. The figure shows that units with extensions had a less pronounced drilling spike prior to the expiration of the original primary term and a larger drilling spike prior to the expiration of

Figure 4: Time series of the Henry Hub, Louisiana natural gas 12-month futures price, Haynesville leases signed, and the number of Haynesville wells spudded

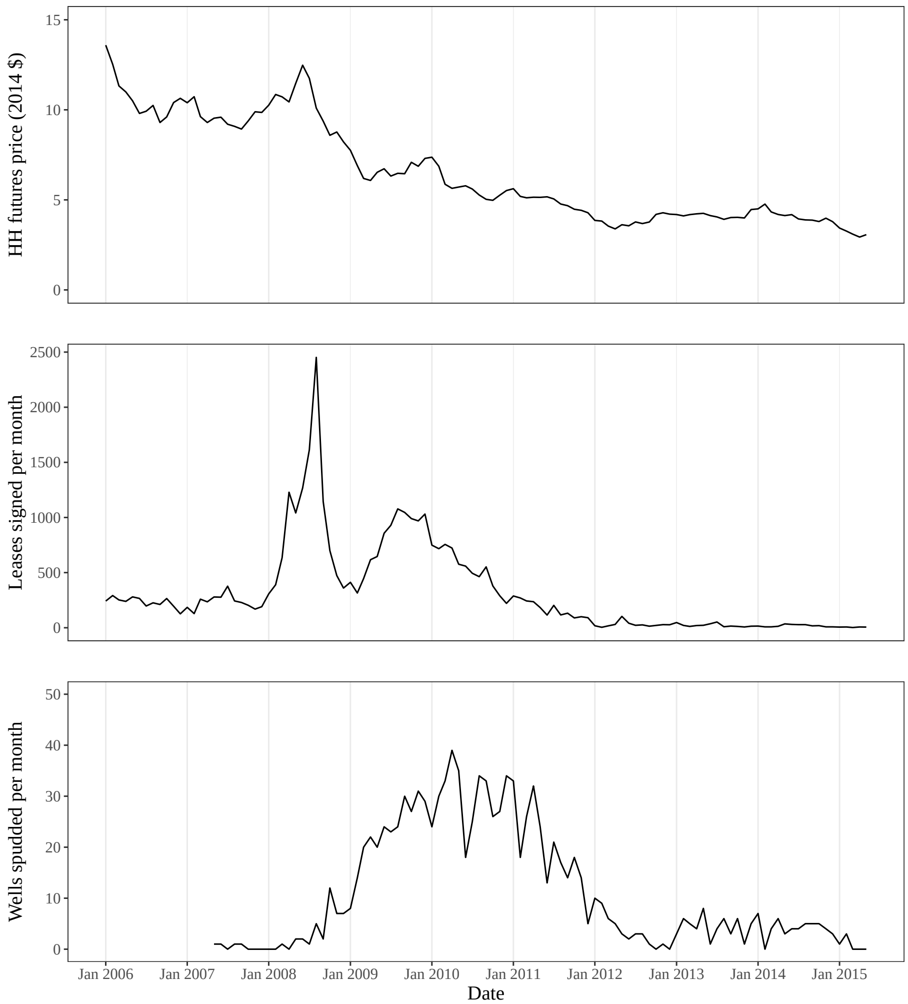

Note: Data include activity for units in our analysis sample, as defined in Section 3.4.

the extension term two years later.

It may be rational for operators to drill after the first lease expires if that lease does not account for a large share of the overall leased acreage and the remaining leases do not expire for some time. To examine this possibility, we calculate the amount of time from when the first lease expires to the time at which 50% of all acreage would expire and then compare units where that value is above versus below the median. In Panel (b) of Figure 6, we show that when there is a shorter amount of time until 50% of lease acreage expires, there is a more pronounced spike in drilling just before the first lease expires, consistent with this intuition.

Figure 5: Date of first drilling relative to first expiration date

(a) Kernel density **(b)** Histogram

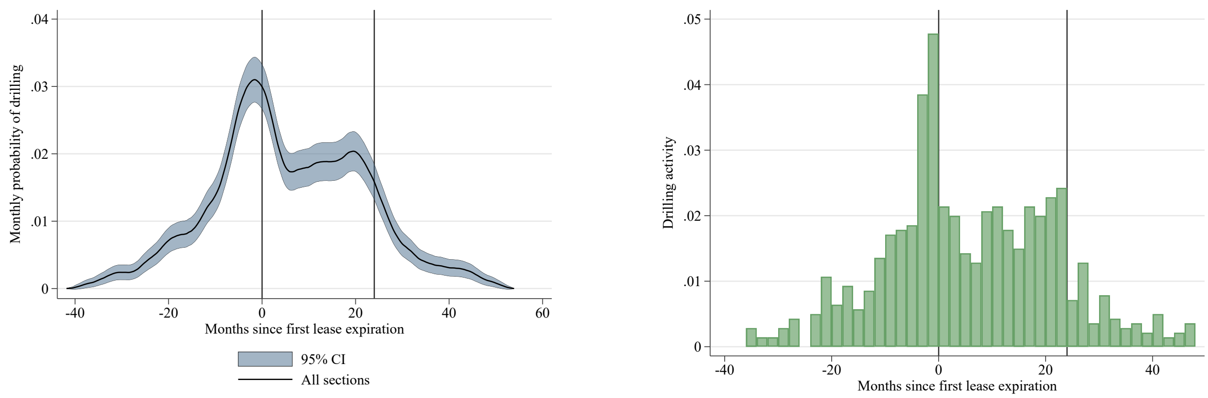

Note: Panel (a) is a kernel-smoothed estimate of the probability of drilling the first Haynesville well in a unit on a given date, relative to the expiration date of the first lease within the unit to expire. Panel (b) is a histogram showing the same data, in which each bar represents two months. Vertical lines are drawn at the date of first lease expiration and two years after first lease expiration.

Figure 6: Timing of first drilling by extension status and number of total wells eventually drilled

(a) Extensions vs no extensions **(b)** Longer vs shorter time until 50% expiration

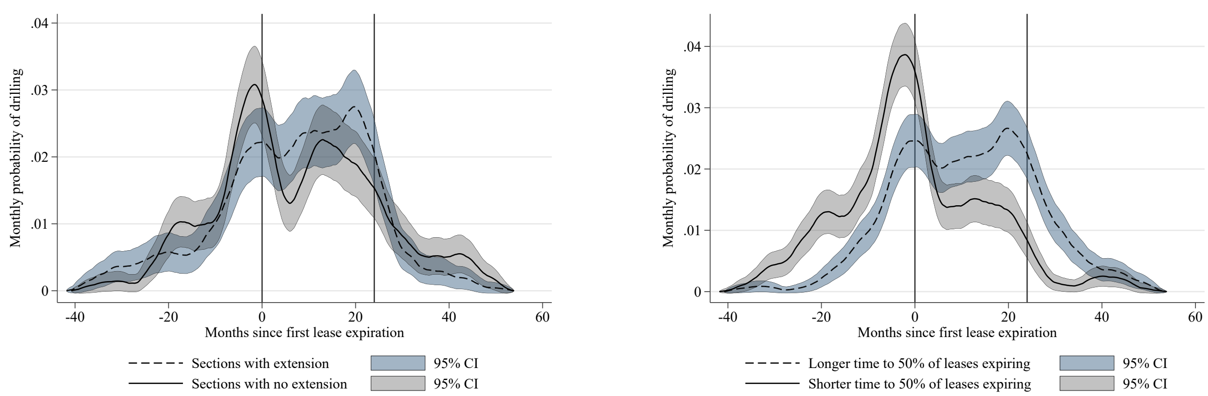

Note: Both panels present kernel-smoothed estimates of the probability of drilling the first Haynesville well in a unit on a given date, relative to the expiration date of the first lease within the unit to expire. Vertical lines are drawn at the date of first lease expiration and two years after first lease expiration. Panel (a) is a comparison of units where the first expiring lease had a built-in two-year extension clause versus units where the first expiring lease did not have any extension clause. Panel (b) compares units where the time elapsed from first lease expiration to expiration of 50% of acreage is longer to those for which it is shorter.

If the primary term is pushing firms to drill a well to hold leased acreage when drilling was otherwise unprofitable, we would expect that many units would have only a single well for an extended period of time. Indeed, of the units in our sample that have drilling, 74% of them have only one well. In Panel (b) of Figure 2, and in Figure 3, we show that a large fraction of drilled units only had one well, even as late as March, 2017.

The variation in shading in Figure 3 represents different unit operators. We find that operators often have control of multiple contiguous units, which suggests that the drilling patterns are not being driven by externalities like common pool inefficiencies or information spillovers (see Kellogg and Reguant (2021) for a review of research on these topics). We examine this possibility further in Figure A.2 in Appendix B.2, where we show that there is a spike in drilling prior to primary term expiration regardless of whether the unit operator controls nearby units or not.

5 Specification and estimation of a model of firms' drilling and input choices

While the results presented in Section 4 show that primary terms cause bunching in drilling timing, they do not tell us what would have happened in the absence of primary term deadlines. For instance, the analysis does not tell us when wells drilled just prior to expiration would have been drilled in a "no primary term" counterfactual nor whether they would have been drilled at all. Moreover, if firms are forward-looking, the primary term can also hasten drilling substantially before the deadline. The bunching analysis also does not shed light on the effects of the royalty nor on any impacts of lease terms on water input choices.

To simulate these counterfactuals and then evaluate the effects of primary terms on surplus outcomes, we develop a model of firms' drilling timing and water input choice problem. This section discusses how we specify and estimate: (1) the well-level production function, including the effect of water use on wells' output; (2) well-level profits; (3) the time series processes for the prices of natural gas and drilling inputs; and (4) firms' dynamic optimization problem. We provide a summary of all estimated parameters near the end of this section in Table V. Additional detail on the model and its estimation is provided in Appendix C.

5.1 Production function

We specify the production function as $Y_{ij} = g(\theta_i, X_i, W_j, \varepsilon_{ij})$ for a well j drilled in unit i. Y_{ij} denotes the present value of well j's cumulative lifetime gas output (in millions of mmBtu, henceforth trillions of Btu (TBtu)), θ_i are factors observed by the firm but not by the econometrician, X_i represents observable covariates that may affect productivity, W_j is water input (in gallons), and ε_{ij} represents post-drilling output shocks that, unlike X_i and θ_i, are unknown to the firm at the time of drilling.

We include water inputs W_j in our production function because we are interested in capturing how lease royalties may reduce firms' use of costly well inputs, beyond just distorting drilling timing alone. These input distortions may be especially important in the modern shale industry, in which water is used to fracture the hydrocarbon-bearing rock formation and to convey proppant (specialized sand) that keeps fractures open. We focus on water in particular because it is the input most commonly recorded in our data; the use of proppant and other inputs is infrequently reported. We therefore think of measured water input W_j as a proxy for all inputs into hydraulic fracturing; i.e., not just water itself but also proppant (which Covert (2015) finds is correlated with water), labor inputs, pumping equipment, chemical additives, and so on. Similarly, the water price series P_{wt} that we estimate in Section 5.3 should be interpreted as capturing the marginal cost of this collection of inputs.

To estimate the production function, we start with the full sample of wells that have both production and water records (see table I) and that can be mapped to a Haynesville unit.[7] We limit the sample to those wells that were completed between 2009 and 2013, when most of the Haynesville drilling was completed and the state of fracking technology used in the Haynesville was stable (Cadotte et al. 2017). For observable geology X_i, we use original gas in place (OGIP_i). We then specify the production function as equation (1):

$$Y_{ij} = \beta_0 + \beta_1 X_i + \theta_i + \beta_w \log W_j + \varepsilon_{ij} \tag{1}$$

Water input W_j enters the production function in log form to reflect decreasing marginal returns to water. This production function is additively separable with output Y_{ij} measured in levels, not logs as in Covert (2015) or other applications using Cobb-Douglas assumptions. We use this functional form because we find that water use and OGIP are, if anything, slightly

[7]Because wells vary slightly in horizontal length, we scale both water inputs and cumulative production by well horizontal length such that both are standardized for a well with a length of 1485 meters.

negatively correlated, rather than positively correlated as would be the case were $\log Y_{ij}$ on the left hand side of Equation (1).[8] Unlike Covert (2015), which models spatial heterogeneity in the marginal productivity of water because that paper is primarily interested in how firms learn about their local production function, we simplify the production function by assuming there is a constant β_w over space.

Estimating Equation (1) using OLS would yield inconsistent estimates for two reasons. The first is a standard selection problem: if the decision to drill is selected on the unobservable θ_i, then β_1 will be biased towards zero per the logic in Olley and Pakes (1996). The second is measurement error that will attenuate the estimate of β_w.[9] One source of measurement error is incorrect recording of the volume of water used. For instance, we find that where our data overlap with data in the FracFocus repository, wells' water use records do not always match. A second and likely more important source of measurement error is a difference between the total volume of water used and that which actually contributes to gas production — only the latter of which is represented by the $\log W_j$ term in equation (1). Frac jobs typically involve substantial losses of water into the formation that do not contribute to the generation of fractures or, ultimately, gas production. These losses can vary substantially across wells and are difficult for the firm to predict in advance because they depend on the presence of natural fractures near the wellbore and the properties of the frac fluid as it travels through the well and formation (Penny et al. 1985; Montgomery 2013; Yarushina et al. 2013). As a result, recorded water use exhibits a remarkably large variance—with many outliers—that is largely unexplained by location or time effects.[10] Under an assumption that these water losses (in logs) are independent of $\log W_j$, they manifest in our model as classical measurement error.

We therefore use a two-step procedure to identify all of the parameters in the production function given in Equation (1). We first estimate the coefficient on $\log W_j$, β_w, using a semi-parametric instrumental variable specification. We later estimate β_0 and β_1, along with

[8] As OGIP is a significant driver of output and drilling decisions, a Cobb-Douglas production function would imply that W_j should increase with OGIP_i if firms choose W_j optimally. Instead, the correlation between W_j and OGIP_i in our data is -0.127 (-0.114 after controlling for month fixed effects). Because this correlation is modest, and because specification (1) greatly simplifies the model's solution (since each well's optimal W_j will be independent of its OGIP_i and θ_i), we model production per specification (1) rather than use a specification that allows the marginal productivity of water to decline with OGIP_i and θ_i.

[9] The usual simultaneity concern in production function estimation—that unobserved productivity heterogeneity would bias the estimate of β_w upwards—is not present here due to the additively separable functional form of Equation (1).

[10] We plot the dispersion of water use in Figure A.4 in Appendix B.2.

the variance of θ_i, by applying a maximum likelihood estimator to our full model, which explicitly accounts for selection of which units are drilled (see Section 5.4).

Our strategy for identifying β_w relies on the fact that firms' choice of water input W_j should depend on both the price of natural gas P_t and the price of water P_{wt}. Because we do not observe P_{wt}, we use an instrumental variable approach where we predict water use for each well using water use from all other wells drilled within the same month while partialling out the contribution of geology via non-parametric controls for latitude and longitude (Robinson 1988). We do this using the IJIVE estimator proposed by Ackerberg and Devereux (2009). Further details are in Appendix C.1. The identifying assumptions for IJIVE are that there is no correlation in measurement error across wells drilled in the same month, and that there are no time-varying factors other than P_t and P_{wt} that would affect Haynesville-wide time series variation in water use and production during our sample. A potential violation of this second assumption would be factor-augmenting technological progress, but our understanding is that significant improvements in technology occurred after but not within our 2009–2013 sample period (Cadotte et al. 2017). In fact, during this period both water use and production from new wells actually slightly decreased on average (consistent with an overall decrease in P_t and an increase in completion costs) rather than increased (see Figure A.4 in Appendix B.2.)

Using IJIVE, we find an estimate of β_w equal to 2.41 TBtu. This estimate implies that moving from the 25th to the 75th percentile of water (from 4.4 million gallons to 6.6 million gallons) would increase production by 0.97 TBtu, which is approximately equal to one-half the interquartile range of production (1.85 TBtu). The estimates are presented in Table IV, column 1. We obtain a similar estimate of β_w equal to 2.53 TBtu (column 2) using the UJIVE methodology of Kolesàr (2013), which takes a similar leave-one-out instrumental variable approach but handles covariates differently (see appendix C.1). Column 3 shows the 2SLS estimate where we use month fixed effects as an instrument and include flexible controls for latitude and longitude in both stages (Li and Stengos 1996). Column 4 shows the OLS estimate where we flexibly control for latitude and longitude (Robinson 1988). Consistent with classical measurement error and the discussion in Ackerberg and Devereux (2009), both the 2SLS and OLS estimates are biased towards zero, with the bias being greater for OLS than for 2SLS.

The IJIVE and UJIVE estimates will themselves be biased towards zero if measurement error is correlated across wells drilled in the same month. This assumption is not directly

Table IV: Estimates of the impact of water input on gas production

	IJIVE	UJIVE	2SLS	OLS
β_w	2.41	2.53	1.86	1.15
	(0.89)	(0.95)	(0.45)	(0.22)
	[0.76]	–	[0.61]	[0.22]
N	2,019	2,019	2,019	2,019

Note: Estimates of β_w from Equation (1). Production is measured in millions of mmBtu (TBtu); water is measured in millions of gallons. Standard errors are clustered at the township level, with analytic standard errors in parentheses and bootstrapped standard errors in brackets (using 5,000 bootstrap draws). Due to UJIVE's computational burden, we do not compute bootstrapped standard errors for UJIVE. The IJIVE and UJIVE estimators do not produce traditional first-stage F statistics, so to assess first-stage fit we instead project observed residualized log water onto predicted residualized log water. For IJIVE, the coefficient from that projection is 0.52 (standard error = 0.18), and for UJIVE it is 0.50 (standard error = 0.17).

testable, though we have found that if we apply the IJIVE estimator while leaving out all wells drilled in the same section, we obtain a larger but imprecise estimate of β_w equal to 3.37 TBtu (clustered standard error = 2.14 TBtu). We have therefore also simulated our main counterfactuals discussed in Sections 6 and 7 below using a value of $\beta_w = 3.50$ TBtu, finding results that are qualitatively similar.[11]

5.2 Drilling profits

If the firm operating in unit i drills its first well j in period t, it gets well-level profits π_{ijt} that depend on profit-maximizing water use W^*_{ijt}, the production function $g(\theta_i, X_i, W^*_{ijt}, \varepsilon_{ij})$, the natural gas price P_t, the royalty rate k_i, the rig dayrate D_t, operating and gathering costs c, and the share of acreage remaining under lease at the time of drilling f_{it}. Profits also depend on the severance tax s, the corporate income tax τ, and the effective corporate income tax

[11] When we set $\beta_w = 3.50$ TBtu, we also re-estimate the water price intercept γ_0 in Equation (5) and the β_0, σ_ε, and α_0 parameters estimated via maximum likelihood in Section 5.4. With these parameters, the analysis in Section 6 continues to find modest effects of a primary term on total surplus in the presence of a 25% royalty, with the effect positive for units with production greater than 2.4 TBtu. This result is consistent with the intuition discussed in that section that primary terms are less effective at increasing surplus when output depends strongly on firms' input choices. The analysis from Section 7 continues to find that a primary term slightly reduces the landowner's revenue.

rate on capital expenditure τ_c. The functional form for profits is given by Equation (2):

$$\begin{aligned}\pi_{ijt} &= f_{it}(1-\tau)((1-s)(1-k_i)-c)P_t g(\theta_i, X_i, W^*_{ijt}, \varepsilon_{it}) \\ &\quad - (f_{it}(1-\tau)(1-s+sk_i)+\tau-\tau_c)(\alpha_0+\alpha_1 D_t) \\ &\quad - f_{it}(1-\tau)(1-s+sk_i)P_{wt}W^*_{ijt}. \end{aligned} \tag{2}$$

The royalty k_i is the acreage-weighted average royalty for the unit. The fraction of the unit under lease f_{it} determines what fraction of total profits go to the firm, as mineral owners whose leases have expired become minority shareholders in the lease.

At the time of drilling, the firm pays a total rig rental payment that depends on the rig dayrate D_t. We set α_1 to be the average number of days required to drill a Haynesville well in our sample: 59.3. The intercept α_0 is unobservable fixed costs and is estimated via maximum likelihood as discussed in subsection 5.4. Costs of water and associated fracking inputs are denoted $P_{wt}W^*_{ijt}$, where P_{wt} is the price of water (the first-order condition for W^*_{ijt} is given in Appendix C.2).

The term c denotes operating and gathering costs paid to gathering pipeline operators and other service providers. We treat these costs as proportional to the value of the gas produced. To calibrate c, we use Gülen et al. (2015), which states that typical operating and gathering costs cP_t for the Haynesville shale were about \$0.60/mmBtu in the earliest years of Haynesville shale gas extraction. We divide \$0.60/mmBtu by the average natural gas price prevailing in 2009–2010 to arrive at $c = 0.0963$.

The severance tax s in Louisiana's shale wells is 4% (Kaiser 2012).[12] The combined state and federal marginal corporate income tax rate τ is 40.2%, and the corporate income tax on capital drilling expenditure τ_c is 36.8% (Metcalf 2010; Gülen et al. 2015).[13] We assume that both s and τ also apply to the owner's royalty income.

[12] Louisiana's severance tax on Haynesville shale wells becomes payable after either the well has been producing for two years or the well's drilling costs have been paid, whichever comes first (Kaiser 2012). We simplify by assuming a tax of 4% on production revenue and allowing the firm to deduct drilling costs (subject to revenue exceeding costs). In the event that the well does not "pay out", some taxes are not owed and the firm's profit is given in Equation (A.7) in Appendix C.2.

[13] To calculate τ_c, we treat 50% of drilling expenditures as immediately expensable, while the remainder must be capitalized and depreciated over time using the double declining balance method (Metcalf 2010).

5.3 Price process and water price estimation:

When firms decide when to drill a well, they must form expectations about future prices. Following Kellogg (2014), we model natural gas prices P_t and rig dayrates D_t as following the Markov processes given by equations (3) and (4), respectively. The drift parameters κ_0^P, κ_1^P, κ_0^D, and κ_1^D 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 2008. 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.[14] We treat each period t as a calendar quarter and aggregate all data to this level.

$$\ln P_{t+1} = \ln P_t + \kappa_0^P + \kappa_1^P P_t + \sigma^P \eta_{t+1}^P, \tag{3}$$

$$\ln D_{t+1} = \ln D_t + \kappa_0^D + \kappa_1^D D_t + \sigma^D \eta_{t+1}^D. \tag{4}$$

Parameter estimates are shown in Table V and imply that the long-run mean natural gas price is \$3.92/mmBtu and the long-run mean dayrate is \$7258 per day. We assume that the shocks η_{t+1}^P and η_{t+1}^D are drawn from an i.i.d. bivariate standard normal distribution, with a covariance matrix that we estimate using the residuals of equations (3) and (4).

We do not observe the water price P_{wt} directly. We therefore model P_{wt} as a function of P_t, D_t, and unobserved factors ω_t per Equation (5):

$$\log P_{wt} = \gamma_0 + \gamma_1 \log P_t + \gamma_2 \log D_t + \omega_t. \tag{5}$$

To estimate Equation (5), we combine it with the first-order condition for water implied by Equation (2) to obtain a projection of $\log W_j$ onto $\log P_t$ and $\log D_t$ that recovers γ_0, γ_1, and γ_2. We obtain γ_0 = -3.59, γ_1 = 1.01, and γ_2 = 0.24 (see Table V, where the standard errors account for the variance in the estimate of β_w). We apply Bayesian shrinkage to the residuals to back out the disturbance ω_t, and hence log P_{wt}, in each period. Further details are in Appendix C.3.

The estimated ω_t have an AR1 coefficient of 0.09 (with a standard error of 0.24). We therefore assume that the ω_t are not serially correlated, and we also assume that the ω_t are not realized until after firms make their drilling decisions each period. These tractability

[14] $\bar{P}_t$ and $\bar{D}_t$ denote the average price and dayrate, respectively, over 1993–2008.

assumptions imply that while ω_t affects firms' water choice and production for any drilled wells, the dynamic model need not track ω_t as a state variable.

5.4 Dynamic problem and simulated maximum likelihood

We now turn to the firm's dynamic drilling problem. The firm solves an optimal stopping problem, deciding in each quarter whether to drill or to wait. For each unit i in each period t, the firm faces the following state variables when making its decision:

- The gas price P_t and rig dayrate D_t
- The unit's royalty k_i and the unit's schedule of acreage expiration, given by the vector $\mathbf{f_{it}}$. The elements of $\mathbf{f_{it}}$ are f_{it}, $f_{i,t+1}$, $f_{i,t+2}$,..., which denote the share of unit acreage that will remain leased (i.e., unexpired) in t and all future periods.
- Observable (to the econometrician) productivity shifters X_i (i.e., OGIP), and an unobservable (to the econometrician) productivity shifter θ_i

When the firm drills its initial well in some period t, it can either drill one or multiple wells at that date. Drilling one well gives it an infinite timeline to drill additional wells with the acreage held at f_{it}. In our estimation and counterfactuals, we assume that the pooling unit has space for a total of $M = 3$ wells.[15] Conditional on no previous drilling, the firm can choose to drill $m = 0$, 1, 2, or 3 wells. All wells drilled within the unit are assumed to have the same expected productivity ($\beta_0 + \beta_1 X_i + \theta_i$); we abstract away from learning or other interactions between wells drilled in the same unit that may cause optimal water use or production to vary across wells.

Payoffs to drilling and not drilling each potential well also involve cost shocks ν_{it}^1 and ν_{it}^0, respectively, which are i.i.d. draws from a type-1 extreme value distribution with standard deviation σ_ν. These shocks are necessary for rationalizing the data because our model abstracts away from factors such as rig availability, well interference, learning, and financial frictions that might affect firms' drilling timing decisions (Hodgson 2018; Steck 2018; Agerton 2020; Gilje et al. 2020). The larger σ_ν is, the less sensitive are the model's simulated drilling probabilities to factors such as prices, unit productivity, and impending lease expiration.

[15] The mean number of wells in a unit, conditional on at least 2 wells being drilled, is 3.2.

Let S_{it} denote the set of observable state variables (all but θ_i). The Bellman equation is:

$$V_{i,t}(S_{it},\theta_i) = E_{P_{wt},\nu^1_{it},\nu^0_{it}}\Bigg[\max\Bigg\{\delta E[V_{i,t+1}|S_{it},\theta_i] + f_{it}M\nu^0_{it},$$
$$\max_{m\in\{1,..,M\}}\Big[m(E_{P_{wt}}[\pi_{it}(S_{it},\theta_i)] + f_{it}\nu^1_{it}) + (M-m)(\delta E[\tilde{U}_{i,t+1,t}|S_{it},\theta_i] + f_{it}\nu^0_{it})\Big]\Bigg\}\Bigg] \quad (6)$$

$V_{i,t}(S_{it},\theta_i)$ gives the firm's unit-level expected value at date t after learning S_{it} and θ_i, but prior to learning P_{wt}, ν^1_{it}, and ν^0_{it}. δ is the discount factor, set equal to 0.909 (Kellogg 2014). The first line of Equation (6) is the continuation value of not drilling. In the second line, $E_{P_{wt}}[\pi_{it}(S_{it},\theta_i)]$ is the expectation of static well-level profits $\pi_{it}(S_{it},\theta_i)$ from Equation (2), taken over water prices P_{wt}.[16] The total number of wells the firm chooses to drill when it first drills is m, and $E[\tilde{U}_{i,t+1,t}|S_{it},\theta_i]$ gives the per-well continuation value of a unit that is held by production as of date t. In Appendix C.4, we show that this model implies an ordered logit specification for drilling 0, 1, or M wells each period.

The parameters to be estimated are: β_0 and β_1 from the production function (Equation (1)), the drilling cost intercept α_0, the scale parameter σ_ν of the ν_{it} cost shocks, the standard deviation σ_θ of the θ_i, and the standard deviation σ_ε of the ε_{ij}. We assume θ_i and ε_{ij} have mean-zero lognormal distributions. We estimate this set of parameters—which we denote by Ω—using simulated maximum likelihood, similar to Kellogg (2014) and Agerton (2020). This nested fixed point procedure (Rust 1987) simulates, for given parameters Ω, unit-level probabilities that the unit's first well is drilled in each period t and expected output $E_{\varepsilon_{ij}}[Y_{it}]$ conditional on drilling, given the unit's time path of observed state variables $\{S_{it}\}$ (which we denote $\mathbf{S_i}$), a value for θ_i, and the time path of water prices $\{P_{wt}\}$. This empirical strategy assumes that units' royalties k_i and schedules of acreage expiration $\mathbf{f_{it}}$ are independent of θ_i. This assumption is a consequence of the fact that we do not otherwise impose structure on how lease terms are determined in our econometric model. In practice, the observed variation in lease terms is limited but exhibits some dependence on units' OGIP, as discussed in Section 3.3 and Appendix B.2.

For each unit i, let the indicator variable I_{it} equal 1 if unit i is first drilled in quarter t in the data. The indicator I_{i0} equals 1 if the unit is not drilled by Q4 2013, so that

[16] In practice, we have found that $E_{P_{wt}}[\pi_{it}(S_{it},\theta_i)]$ is nearly indistinguishable from $\pi_{it}(S_{it},\theta_i)$ evaluated at $E[P_{wt}|S_{it}]$, so to save computation time our model implements the latter when computing drilling probabilities and continuation values.

$\sum_t I_{it} + I_{i0} = 1\ \forall i$. Let $Pr(I_{it} = 1|\theta_i, \Omega, \mathbf{S_i})$ be the model's simulated probability that unit i is drilled in quarter t, with $Pr(I_{i0}|\theta_i, \Omega, \mathbf{S_i})$ denoting the simulated probability that the unit is not drilled. For drilled units, let $f(Y_{ij}|\theta_i, \Omega, S_{it}, P_{wt}, t)$ denote the pdf of the well's simulated production conditional on drilling at t, evaluated at the well's actual production Y_{ij}. The simulated likelihood integrates over the distribution of θ_i, denoted as $\psi(\theta_i|\sigma_\theta)$ (see Appendix C.6.1 for details). The log likelihood, letting $\{\mathbf{S_i}\}$ denote the set of $\mathbf{S_i}$, is:

$$LL(\Omega, \{\mathbf{S_i}\}, \{P_{wt}\}) = \sum_i \left[I_{i0} \cdot \log \left(\int_\theta Pr(I_{i0} = 1|\theta, \Omega, \mathbf{S_i}) \psi(\theta|\sigma_\theta) d\theta \right) \right]$$
$$+ \sum_i \sum_t \left[I_{it} \cdot \log \left(\int_\theta Pr(I_{it} = 1|\theta, \Omega, \mathbf{S_i}) f(Y_{ij}|\theta, \Omega, S_{it}, P_{wt}, t) \psi(\theta|\sigma_\theta) d\theta \right) \right] \quad (7)$$

We estimate the model using a subset of our analysis sample that we defined in Section 3.4 and used in Section 4. For each unit, define its "start date"—after which time the model considers drilling to be possible—as the date at which leased acreage in the unit reaches its maximum. We restrict the sample to units: (1) whose start date is within the 2009–2013 period used to estimate β_w (as discussed in subsection 5.1); (2) whose leased acreage weakly declines over time after the start date;[17] and (3) that do not experience drilling before the start date. We impose these restrictions to remain agnostic about the process of adding acreage to an existing unit and to avoid modeling how units starting before 2009 survive without drilling into the 2009–2013 sample window. We also drop units for which royalty rates or leased acreage data are either missing or likely to be inaccurate (see Appendix C.5). Our final *estimation sample* contains 241 units, of which 73 experienced drilling. This rate of drilling is lower than that of analysis sample from Section 3.4 (in which 712 of 1,226 units were drilled). This difference is likely a consequence of the fact that the estimation sample selects units that reached their start date in 2009 or later.

5.5 Estimates, model fit, and drilling profits

The maximum likelihood estimates of $\Omega \equiv \{\beta_0, \beta_1, \sigma_\theta, \sigma_\varepsilon, \alpha_0, \sigma_\nu\}$ are presented in Table V, where the standard errors are clustered on township (to account for spatial correlation across nearby sections) and account for the fact that the parameters associated with water

[17]Following the evidence from Section 4, we assume that leases expire at the end of the potential extension term if we observe an extension clause. We assume no further extensions are possible after the initial extension.

Table V: Summary of the model's parameter estimates

Parameter	Notation	Value	Std. error	Units	See section
Taxes and operating / gathering costs					
Severance tax rate	s	4%			5.2
Federal and state income tax rate	τ	40.2%			5.2
Effective income tax on drilling expenditure	τ_c	36.8%			5.2
Operating and gathering cost parameter	c	0.10			5.2
Gas price and rig dayrate transitions					
Price drift constant	κ_0^P	0.0044	(0.0069)		5.3
Price drift linear term	κ_1^P	-0.0024	(0.0013)	mmBtu/$	5.3
Price volatility	σ^P	0.10	(0.018)		5.3
Dayrate drift constant	κ_0^D	0.0044	(0.0069)		5.3
Dayrate drift linear term	κ_1^D	-1.15	(0.64)	$million^{-1}	5.3
Dayrate volatility	σ^D	0.089	(0.016)		5.3
Price - dayrate correlation	ρ	0.33	(0.12)		5.3
Water use and pricing parameters					
Production function coef. on log water input	β_w	2.41	(0.76)	TBtu	5.1
Log water price projection constant	γ_0	-3.59	(4.50)		5.3
Log water price projection coef. on log P_t	γ_1	1.01	(0.18)		5.3
Log water price projection coef. on log D_t	γ_2	0.24	(0.19)		5.3
Well productivity					
Constant	β_0	-37.2	(0.8)	TBtu	5.4
Coefficient on OGIP	β_1	17.3	(4.4)	TBtu/OGIP	5.4
Standard deviation of log(θ)	σ_θ	0.79	(0.2)		5.4
Standard deviation of log(ε)	σ_ε	0.70	(0.16)		5.4
Drilling and completion costs					
Drilling cost dayrate coefficient	α_1	59.3	(1.2)	days	5.2
Drilling cost intercept	α_0	3.26	(2.34)	$million	5.4
Cost shock scale parameter	σ_ν	1.42	(0.78)	$million	5.4

Note: TBtu refers to trillions of Btu (i.e., millions of mmBtu). OGIP is measured in trillions of cubic feet of gas per square mile. In the production function, water enters in gallons. In the log water price projection, the water price is measured in $ per million gallons, the gas price is in $ per TBtu, and the dayrate is in $ per day. The estimated values for α_0 and σ_ν are pre-tax.

Figure 7: Fit of estimated model to the drilling data

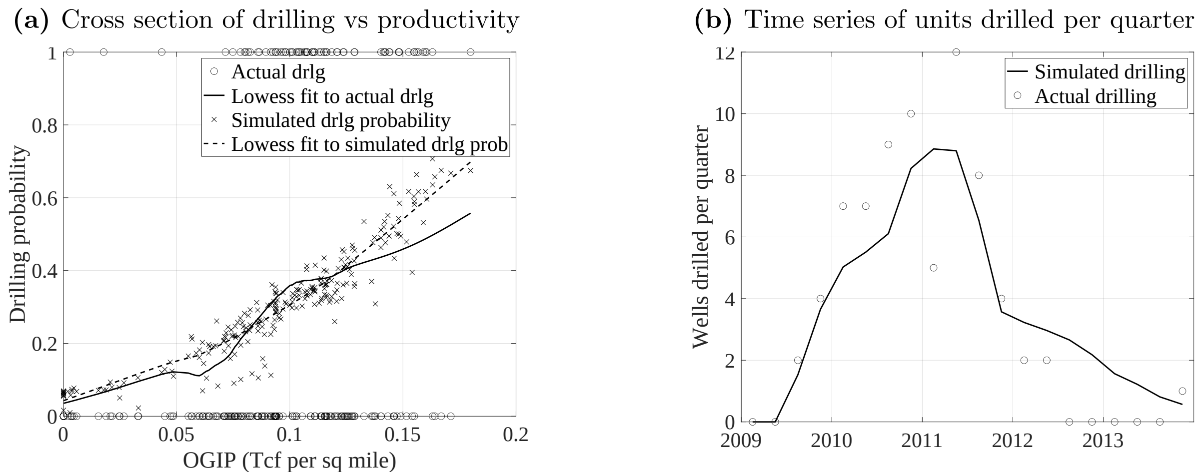

Note: In Panel (a), we plot the following unit-level variables against each unit's OGIP: whether the unit was ever actually drilled (circles plotted at 0 or 1), a lowess fit to actual drilling (solid line), the simulated probability the unit is ever drilled (x's), and a lowess fit to the simulated probability (dashed line). In Panel (b), we plot the number of times each quarter in which a unit is drilled for the first time, in both the actual data (circles) and estimated simulation (solid line). Plotted data include all units in the estimation sample.

input choices ($\beta_w, \gamma_0, \gamma_1$, and γ_2) and rig rental costs (α_1) are estimated in earlier stages (see Appendix C.6.2). In Figure 7, we show that the unit-level drilling simulated by the estimated model fits the data well in both the cross-section (Panel (a)) and time series (Panel (b)).

The estimated scale parameter σ_ν of the type-1 extreme value drilling cost shocks is \$1.42 million, before taxes. The estimate of the drilling cost constant α_0 is pinned down by fitting the model's overall simulated rate of drilling (73.0 of 241 units at the parameter estimates) to actual drilling (73 of 241 units). The estimate of α_0 = \$3.26 million implies that the estimated average drilling and completion cost (including water input) for a well drilled during our 2009–2013 sample is \$10.1 million, which is in the range of accounting costs summarized in Table I.[18]

We find that OGIP and unobservable productivity are empirically important features of the production function. For the output Equation (1), we estimate the coefficient on OGIP $\beta_1 = 17.3$ TBtu per unit of OGIP (measured in trillion cubic feet per square mile), and $\sigma_\theta = 0.789$ (implying a standard deviation of θ in levels of 1.27 TBtu), both of which are

[18]The \$10.1 million estimate includes rig expenditures evaluated at the sample average dayrate, average water expenditures based on water prices during the sample and simulated water input per well, and the expected realization of the cost shock conditional on drilling.

Figure 8: Estimated firms' profits from the first well drilled in each unit of the estimation sample versus the percentage of unit acreage scheduled to expire

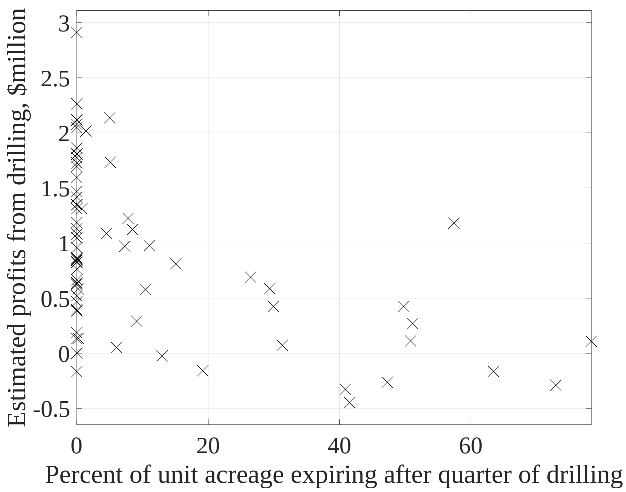

Note: Each observation represents the first well drilled in each of the 73 units drilled in the estimation sample. The horizontal axis plots the percentage of acreage in the unit that is scheduled to expire following the quarter-of-sample in which drilling occurred. The vertical axis plots each well's estimated profits.

statistically significant at the 1% level. The estimate of β_1 is driven by the correlation of both output and drilling hazards with OGIP in the data. Had we not modeled the unobservable θ, the estimate would be biased towards zero due to selection of drilling on θ (Olley and Pakes 1996). Indeed, the estimate of β_1 obtained from simply projecting well-level output on unit-level OGIP in the estimation sample is just 12.7 TBtu per unit of OGIP. Moreover, when we project output on OGIP and on the model's simulated probabilities (at the parameter estimates) that each unit is ever drilled in-sample, the coefficient on the drilling probability is -6.3 TBtu. Consistent with selection on θ, this negative coefficient implies that wells drilled at locations and times when the model thinks drilling would be unlikely, based on all observables S_{it}, are more productive than expected based on OGIP alone. This correlation between output (conditional on OGIP) and drilling probability drives the estimate that σ_θ is large and statistically significant.

We use the estimated model to quantify the realized profits firms earned from the first well drilled in each unit in the estimation sample.[19] In Figure 8, we show how these profits vary with the share of unit acreage that is scheduled to expire within one calendar quarter

[19]Our estimates of realized profits account for the expected value of the cost shock ν, conditional on drilling. We also integrate over the distribution of productivities θ for each unit, using Bayes' rule to infer the probability that a firm of productivity θ drilled the unit.

of drilling. Wells drilled when no acreage is about to expire (the observations on the vertical axis) tend to be highly profitable, earning \$1.18 million on average. In contrast, wells drilled when more than 10% of acreage is about to expire earn average profits of just \$0.24 million, and 37% of such wells actually lose money. The low or negative profitability of these wells is consistent with the notion that firms drilled them primarily to hold acreage and preserve the option to drill additional wells in the future.

6 Effects of royalties and primary terms on drilling, input choice, production, and total surplus

How do contract terms affect drilling, input choice, production, and surplus? We first use our estimated model to simulate outcomes for a hypothetical unit that has the average OGIP from our estimation sample and $\theta = 0$. We simplify the unit's leasing structure by assuming that it is covered by a single lease. We assume that the unit can accommodate up to 3 wells and that drilling one well is sufficient to hold the unit by production. We fix the unit's start date to be the first quarter of 2009 and simulate outcomes from that date forward, taking an expectation over the possible paths for natural gas prices and rig dayrates.

We study outcomes for four types of contracts: (1) a lease with no royalty and no (i.e., an infinite) primary term, so that the lease maximizes total surplus; (2) a lease with a 25% royalty and no primary term; (3) a lease with no royalty and a 5-year primary term; and (4) a lease with a 25% royalty and 5-year primary term. We choose a 25% royalty because this value is the modal royalty in the data, and we choose a 5-year primary term because the modal lease has a 3-year term and a 2-year built-in extension. When modeling cases involving a primary term, we follow our assumption from Section 5 that the original lessee has a terminal payoff of zero should the lease expire without drilling having taken place. Following an expiration, we capture the remaining option value of the undrilled unit by modeling an infinite-horizon lease with the same royalty and same productivity level.

We find that the royalty and primary term have a large effect on the timing of drilling. In Figure 9, we present quarterly first-well drilling probabilities and hazards for a mean productivity unit under each of these leases. We find that a 25% royalty substantially reduces the likelihood of drilling each period: Absent a primary term deadline, the probability that

Figure 9: First-well drilling probabilities and hazards for a mean-productivity unit, with and without a 25% royalty and 5-year primary term

(a) Quarterly drilling probabilities **(b)** Quarterly drilling hazards

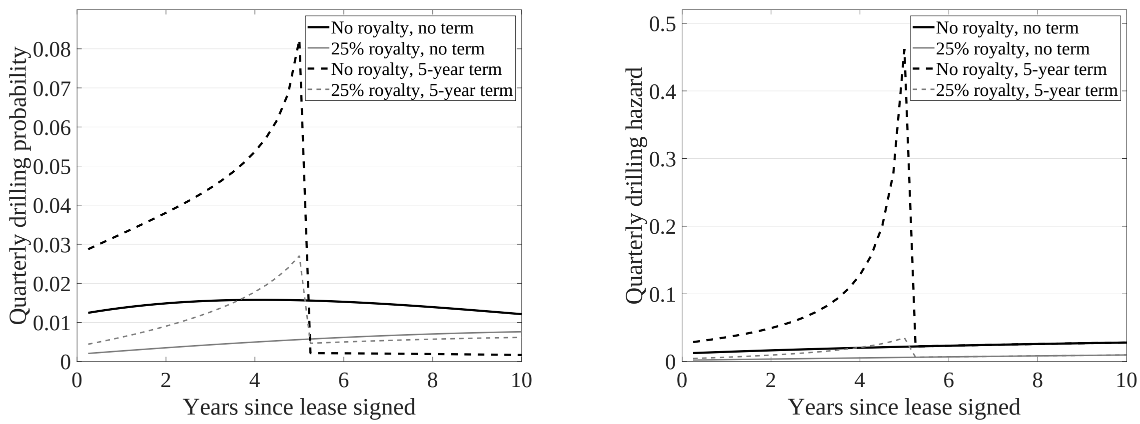

Note: Graphs compare expected drilling timing for four scenarios: 0% royalty and no primary term (black line, and total surplus maximizing), 25% royalty and no primary term (gray line), 0% royalty and 5-year primary term (dashed black line), and 25% royalty and 5-year primary term (dashed gray line). In Panel (a), we show drilling probabilities; in Panel (b), we show drilling hazards. Simulations are for a mean productivity unit (mean OGIP and $\theta = 0$) starting in Q1 2009. Probabilities and hazards shown after primary term expiration correspond to an infinite-horizon lease with the same royalty and productivity. Probabilities and hazards are expectations over all possible price paths starting from conditions in Q1 2009.

the unit is drilled within 10 years is 58% with no royalty and 21% with a 25% royalty.[20] This large effect of the royalty on drilling probabilities is similar to the result in Bhattacharya et al. (2018) that royalties on state-owned parcels in New Mexico substantially reduced drilling there. We further find that, in addition to delaying drilling, the royalty reduces water input and production conditional on drilling. Average simulated water use for a well drilled within 10 years of the unit's start is 7.8 million gallons with no royalty, but it is 5.8 million gallons with a 25% royalty. This input difference leads expected gas production per well to be 2.6 TBtu with no royalty but 1.9 TBtu with a 25% royalty.[21]

Because of the delay and input choice distortions, imposing a 25% royalty on a lease with no primary term deadline reduces total surplus from \$2.60 million to \$1.71 million (these values are post-tax and include expected surplus earned beyond the 10-year window shown

[20]The drilling probability and hazard increase over time in our simulations because the rig dayrate decreases in expectation after Q1 2009.

[21]If we isolate the direct royalty distortion channel (by fixing the price and dayrate at their Q2 2009 levels), we find similar results.

in Figure 9 if drilling does not occur within 10 years). Of that surplus loss, 43% arises from the water input distortion, with the remainder coming from the drilling timing distortion.[22]

In Figure 9, we also show drilling probabilities and hazards under a 5-year primary term, both with and without a 25% royalty. In both cases, the primary term causes drilling to bunch at the deadline, echoing our descriptive results from Section 4. Without the royalty, the primary term causes drilling to be substantially more likely than under the surplus-maximizing counterfactual of no royalty and no primary term. This distortion reduces total surplus from \$2.60 million to \$2.14 million, where the latter number includes the unit's option value after primary term expiration if it is not drilled during the primary term.

When we account for the royalty, however, the increase in drilling probability induced by the primary term draws the drilling rates closer to, rather than farther from, those obtained under the surplus-maximizing lease. During the 5-year primary term (with a 25% royalty), the overall probability of drilling is 25%, relative to 30% under a surplus-maximizing lease and just 8% for a royalty-only lease. Thus, relative to a royalty-only lease, adding a primary term slightly increases total surplus by 2.8%, from \$1.71 million to \$1.76 million.[23]

Productivity varies across the Haynesville Shale, and the effects of primary terms could depend on that productivity. In Table VI, we summarize our simulation results for the mean-productivity unit as well as for a low-productivity unit (10th percentile OGIP and $\theta = 0$), a high-productivity unit (90th percentile OGIP and $\theta = 0$), and for the average over all productivity types implied by our estimates (where the distribution of productivities accounts for both the distribution of OGIP in our estimation sample and the distribution of θ). We plot average (over all productivity types) drilling probabilities and hazards for each lease condition in Figure 10.

The effect of a primary term on drilling and surplus varies depending on the underlying productivity of the unit. Low-productivity units are naturally associated with lower probabilities of drilling and lower output conditional on drilling than are high-productivity units.[24] For each productivity type, the 25% royalty delays drilling and reduces both water use and

[22] We quantify the surplus loss from the water distortion by computing the per-period change in profits and royalties, conditional on drilling, when a royalty is imposed, multiplying these changes by each period's drilling probability in the absence of a royalty, and then taking the discounted sum.

[23] If we were to exclude post-expiration option value from our total surplus calculation, we would find that the primary term decreases total surplus to \$0.76 million.

[24] Low-productivity units have higher water use conditional on drilling than high-productivity units because low-productivity units are more likely to be drilled under high gas prices or low water prices.

Table VI: Simulated impacts of a 25% royalty and 5-year primary term on unit outcomes

Outcome	No royalty no pri term	25% royalty no pri term	No royalty 5-year term	25% royalty 5-year term
Mean-productivity unit				
Prob(drilled within 5 years)	29.7%	7.9%	90.4%	25.3%
Prob(drilled within 10 years)	57.9%	21.5%	94.2%	36.3%
Water use (million gal) | drlg	7.8	5.8	7.5	5.7
Gas production (TBtu) | drlg	2.6	1.9	2.5	1.9
Total surplus ($million)	2.60	1.71	2.14	1.76
10th percentile OGIP unit				
Prob(drilled within 5 years)	1.6%	0.8%	3.6%	1.4%
Prob(drilled within 10 years)	6.9%	3.8%	8.8%	4.4%
Water use (million gal) | drlg	8.3	6.0	8.2	5.9
Gas production (TBtu) | drlg	1.6	0.8	1.6	0.8
Total surplus ($million)	0.37	0.29	0.36	0.29
90th percentile OGIP unit				
Prob(drilled within 5 years)	78.6%	34.3%	100.0%	92.6%
Prob(drilled within 10 years)	95.7%	61.7%	100.0%	95.7%
Water use (million gal) | drlg	7.3	5.5	7.2	5.3
Gas production (TBtu) | drlg	3.4	2.7	3.4	2.6
Total surplus ($million)	8.06	6.04	7.89	6.24
Mean over all productivities				
Prob(drilled within 5 years)	40.1%	22.4%	60.7%	42.2%
Prob(drilled within 10 years)	53.1%	35.0%	64.1%	46.7%
Water use (million gal) | drlg	7.4	5.4	7.4	5.3
Gas production (TBtu) | drlg	3.5	3.0	3.3	2.8
Total surplus ($million)	5.63	4.66	5.49	4.73

Note: Reported drilling probabilities are for the first well in the unit. Reported water use and gas production are for the first well and conditional on drilling within the first 10 years of the unit. The mean productivity, 10th percentile OGIP, and 90th percentile OGIP results assign $\theta = 0$. The "mean over all productivities" results report outcomes that are averaged across the distribution of productivities in the estimation sample of units, accounting for both the distribution of OGIP values in the estimation sample and the estimated distribution $\psi(\theta)$. Reported surplus values are post-tax.

gas production, leading to a substantial reduction in total surplus. Absent a royalty, a 5-year primary term would also reduce total surplus for all productivity types, since it would increase the probability of drilling above that associated with the no-royalty, no primary term contract. But after accounting for the royalty, the increase in drilling induced by the primary term tends to better align the overall likelihood of drilling the unit's first well with

Figure 10: Mean (over all productivities) first-well drilling probabilities and hazards, with and without a 25% royalty and 5-year primary term

(a) Quarterly drilling probabilities

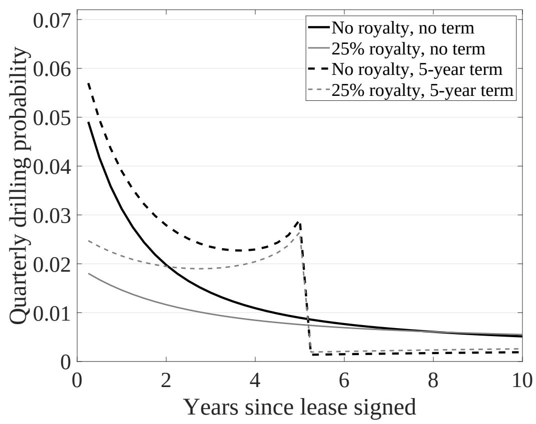

(b) Quarterly drilling hazards

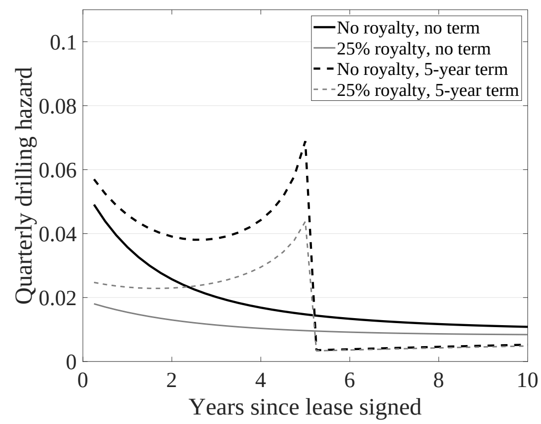

Note: Graphs compare expected drilling timing for four scenarios: 0% royalty and no primary term (black line, and total surplus maximizing), 25% royalty and no primary term (gray line), 0% royalty and 5-year primary term (dashed black line), and 25% royalty and 5-year primary term (dashed gray line). In Panel (a), we show drilling probabilities; in Panel (b), we show drilling hazards. Probabilities are the average over all unit-level productivities (per the distributions of OGIP and θ), starting in the first quarter of 2009. Hazards are computed from the average probabilities and account for selective drilling by higher-productivity types. Probabilities and hazards shown after primary term expiration correspond to an infinite-horizon lease with the same royalty and productivity. Probabilities and hazards are expectations over all possible price paths starting from conditions in Q1 2009.

that achieved by the no-royalty, no primary term contract, as shown in Table VI and Figure 10. The primary term modestly increases total surplus for all but low-productivity units. Averaged across all units, total surplus with a 25% royalty and 5-year primary term is 1.5% greater than total surplus with a 25% royalty and no primary term ($4.73 million vs. $4.66 million). Put another way, the primary term recovers 7.3% of the $0.98 million surplus loss imposed by the 25% royalty.

The increase in total surplus induced by the primary term is modest for three reasons. First, the resulting time profile of drilling probabilities still does not match that of a surplus-maximizing lease—the bunching at the deadline remains distortionary. Second, the primary term only influences the drilling of the first well, not later wells. Third, the primary term does not affect water input choice conditional on drilling. We examine the relative importance of these mechanisms by simulating cases in which we allow only one well per unit or we shut

Table VII: Increase in total surplus induced by a 5-year primary term, as a percent of the surplus loss induced by a 25% royalty

3 wells/unit $\beta_{\mathbf{w}}$ =**2.41 TBtu**	**1 well/unit** $\beta_{\mathbf{w}}$ =**2.41 TBtu**	**3 wells/unit** $\beta_{\mathbf{w}} = \mathbf{0}$ **TBtu**	**1 well/unit** $\beta_{\mathbf{w}} = \mathbf{0}$ **TBtu**
7.3%	21.0%	17.2%	42.6%

Note: Simulated surplus values are post-tax and averaged across the distribution of productivities in the estimation sample of units, accounting for both the distribution of OGIP values in the estimation sample and the estimated distribution $\psi(\theta)$. Values are for a unit starting in Q1 2009 and are taken as an expectation over all possible price paths.

down firms' water input decision by setting $\beta_w = 0$ (and then re-estimating the model's γ_0, β_0, α_0, and σ_ε parameters). We present the results of these simulations in Table VII. When we allow the unit to accommodate only one well, we find that that a 5-year primary term recovers 21.0% of the surplus loss induced by the royalty (averaging across all productivity types). The corresponding percentage for the model with no water input is 17.2%, and that for the model with one-well units and no water input is 42.6%. Thus, primary terms appear to be more effective at counteracting the surplus-reducing effects of the royalty in conventional, non-shale gas settings in which each well can only hold the acreage it drains and large fracking jobs are not required.

7 Mineral owners' revenue-maximizing lease contracts

7.1 Conceptual framework

Thus far, we have studied how royalties and primary terms impact drilling outcomes and total surplus from a mineral lease. This section examines how lease terms affect the expected discounted revenue received by the mineral owner. Addressing this question requires that we model how the up-front bonus payment is determined, which in turn requires a model of the leasing process and mineral owners' objectives and beliefs.

We assume the owner's goal is to maximize expected revenue, that the owner can make a take-it-or-leave-it (TIOLI) lease offer to the firm, and that the firm has private information about productivity.[25] This framework has been used to study effects of royalties in auctions

[25]There may be factors other than expected revenue that enter into mineral owners' objective function.

of state-owned leases, where the royalty reduces the up-front bonus payment while still potentially increasing the owner's total expected revenue (Bhattacharya et al. 2018; Ordin 2019; Kong et al. 2022). As discussed in Hendricks et al. (1993) and Skrzypacz (2013), the royalty compresses firms' type space, reducing their information rent and increasing the owner's total payoff. However, the royalty also distorts firms' drilling incentives, and the owner's revenue-optimal royalty trades off information rent reduction against incentive distortions.

The assumption that the owner can make a TIOLI offer is strong since private mineral leases are typically not allocated using organized auctions (Covert and Sweeney 2019). We adopt it because it allows for a tractable model and for our results to be compared to those from prior work. The fact that Haynesville leases include royalties suggests that owners have at least some bargaining power, since if firms could make TIOLI offers the equilibrium contracts would not include a contingent payment (Skrzypacz 2013).

Unlike papers using oil and gas auction data, we do not observe bonus bids and must make assumptions about what information is common knowledge and what is known only to firms. We assume the owner knows the value of OGIP, the values of all parameters, all price paths, and the distribution $\psi(\theta)$, but that it does not know the true value θ_i.[26] As in our analysis in Section 6, we simplify the unit's leasing structure by assuming it is covered by a single lease, abstracting away from interactions among different mineral owners within a unit.

In our simulations, the owner makes a TIOLI lease offer at $t = 0$ to a single firm of unknown type θ_i.[27] If the firm accepts the offer, then starting at $t = 1$ it operates the lease per the model from sections 5 and 6. The offer includes a bonus, royalty, and primary term (which may be infinite). For a given royalty and primary term, we assume the bonus is set optimally so that it trades off, on the margin, the owner's immediate revenue gain

For instance, owners may be risk-averse or risk-loving, may have discount rates or tax incentives that differ from those of the firm, or may experience disutility from having an oil and gas lease. We have explored this last possibility by simulating an alternative model in which the owner loses \$0.1m of value each year the unit is under contract (and in perpetuity if the unit is drilled). In the specification with one well per unit and no water input, a 5-year primary term increases the owner's value by 6.6% rather than the 0.7% shown in the bottom half of Table VIII.

[26] We also examine a scenario in which the owner knows the distributions of θ and OGIP but not the exact value of either. We find that adding uncertainty about OGIP increases the dispersion of productivity only slightly, and results are similar to those presented here.

[27] A 2010 survey of mineral lessors in the Marcellus Shale in Pennsylvania found that only 21% of them spoke with more than one company before signing a lease (Ward and Kelsey 2011).

Figure 11: Expected owner value and total surplus vs. royalty rate

(a) Endogenous water choice **(b)** Fixed water choice

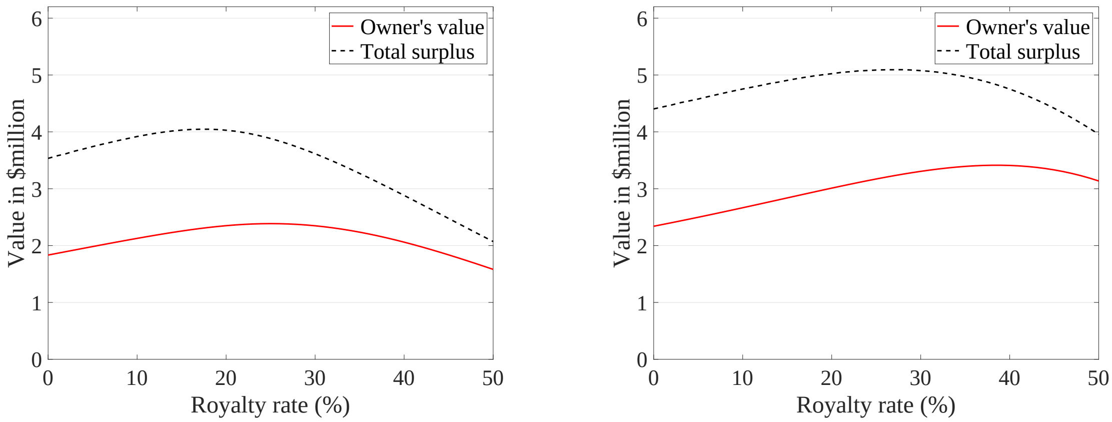

Note: Graphs show how expected owner value (solid lines) and total surplus (dashed lines) change as the royalty rate varies, with no primary term. In Panel (a), we show results when firms endogenously choose water inputs; in Panel (b), we show results with fixed water use. Values are post-tax for a unit with mean OGIP, taken as expectations over $\psi(\theta)$ and all possible price paths starting from conditions in Q1 2009.

from a higher bonus against the loss of revenue from types who decline the offer. Thus, our incorporation of the bonus payment allows for the possibility that the lease may not be executed, unlike in our analysis in Section 6.

7.2 Revenue-maximizing royalties with no primary term

The mineral owner's payoff is maximized by a positive royalty rate. In Figure 11, we show how the royalty rate affects the mineral owner's payoff for a unit with the mean value of OGIP and no primary term. In Panel (a), which presents our baseline specification from Section 5, the owner's revenue-maximizing offer is a royalty rate of 25% with a bonus of $0.86 million. Under these contract terms, there is a 48% probability that the firm accepts the contract. The owner's total expected value is $2.39 million.

When we shut down firms' water input choice (i.e., set $\beta_w = 0$), the owner's revenue-maximizing royalty rate is substantially higher—39% rather than 25%—as shown in Panel (b) of Figure 11. This result highlights that in the shale oil and gas industry, where fracking inputs are an important determinant of production, the revenue-maximizing royalty is affected by moral hazard in firms' input choices, not just drilling timing. Previous work

Figure 12: Expected owner value and total surplus vs. primary term

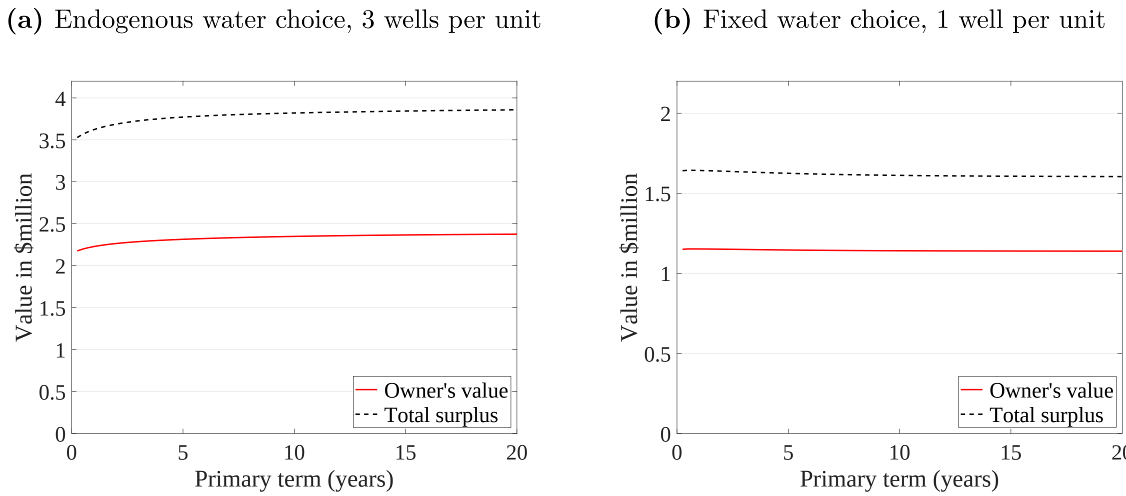

Note: Graphs show how expected mineral owner value (solid lines) and total surplus (dashed lines) change as the primary term varies, holding the royalty fixed. In Panel (a), we show results when firms endogenously choose water inputs, the unit accommodates 3 wells, and the royalty is 25%. In Panel (b), we show results when water use is fixed, the unit accommodates 1 well, and the royalty is 39%. Values are post-tax for a unit with mean OGIP, taken as expectations over $\psi(\theta)$ and all possible price paths starting from conditions in Q1 2009.

on owner-optimal royalties (Bhattacharya et al. 2018; Ordin 2019; Kong et al. 2022) has heretofore only considered the second of these effects.

We also evaluate how the royalty rate affects total surplus, continuing to assume that for a given royalty rate, the bonus is set to maximize the owner's revenue. We find that the surplus-maximizing royalty is lower than the owner's revenue-maximizing royalty, as is clear from the plots in Figure 11. Still, the royalty that maximizes total surplus is strictly greater than zero because a very low royalty rate leads the owner to set a high bonus payment, which excludes a large set of firm types.

7.3 Impacts of primary terms on mineral owners' revenue

To evaluate the impact of a primary term on the owner's revenue, we must model how the owner proceeds should the initial primary term expire without drilling. We assume that upon expiration, the owner makes a TIOLI renewal offer to the firm, wherein the new lease has the same royalty and primary term as the initial lease, and the bonus is revenue-maximizing given the gas price and rig dayrate on the original lease's expiration date. Renewal offers then continue with each lease expiration until the firm either drills or rejects an offer.[28]

We begin by evaluating the owner's payoff at different primary term lengths, using our baseline model evaluated at the sample mean OGIP,[29] and with the owner-optimal royalty of 25% from subsection 7.2. In Panel (a) of Figure 12, we show that the owner's expected value is slightly increasing in primary term length and is in fact maximized at an infinite primary term. In Table VIII, we show the effects of primary terms on equilibrium contracting outcomes and the division of surplus. An infinite primary term results in owner expected revenue of \$2.39 million, while 5-year and 3-year primary terms result in expected revenue of \$2.31 million and \$2.29 million, respectively. This value reduction is associated with a reduction in the initial bonus (\$0.86 million with an infinite term vs. \$0.61 million with a 5-year term) and a reduction in the share of firms that accept the contract (48% with an infinite term vs. 39% with a 5-year term). The reduction in contract acceptance then leads total surplus to increase in primary term length as well, as shown in Figure 12 and Table VIII.

When we shut down water use and one allow the lease to accommodate one well, we find that a primary term can increase surplus. In Panel (b) of Figure 12, we report results from this case, where we set the royalty at the owner-optimal royalty of 39% (per Subsection 7.2).[30] Now, we find that a primary term slightly increases the owner's value, from \$1.14 million with an infinite primary term to \$1.15 million with a 5-year term. As shown in Table VIII, imposing a primary term does not substantially reduce firms' participation under the optimal bonus, and total surplus is then greater with a primary term than without one.

The intuition driving these results stems from two opposing forces. First, as discussed in Section 6, the primary term can increase total surplus by counteracting the royalty's

[28]In modeling renewals, we assume the mineral owner believes it 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 expiration date, and a low-type firm would not have accepted the original contract. However, allowing beliefs to be updated in this manner would result in an intractably complicated state space. Appendix D discusses how we model lease bonuses and renewals in more detail.

[29]The pattern of results shown in Figure 12 and Table VIII—a primary term slightly decreases owner value and total surplus in the baseline model but slightly increases it in the single-well, no water input model—holds for each OGIP decile.

[30]If we instead study this case with a 25% royalty, we find that the primary term slightly decreases the owner's value (from \$1.06 million with no primary term to \$1.04 million with a 5-year term and \$1.03 million with a 3-year term). Primary terms reduce the owner's value in this case because the 25% royalty is too low, such that drilling is under-distorted relative to what would maximize the owner's value.

Table VIII: Summary of simulations of the impacts of royalty and primary term combinations on owner's value and total surplus, with an initial bonus that maximizes owner value

	Owner value	Total surplus	Initial bonus	Share of firms participating
Baseline model				
No royalty, no primary term	\$1.83m	\$3.53m	\$7.84m	23.4%
25% royalty, no primary term	\$2.39m	\$3.88m	\$0.86m	48.3%
39% royalty, no primary term	\$2.10m	\$2.96m	\$0.39m	60.2%
25% royalty, 5 yr. primary term	\$2.31m	\$3.77m	\$0.61m	39.4%
25% royalty, 3 yr. primary term	\$2.29m	\$3.73m	\$0.48m	37.4%
One well per unit, no water input				
No royalty, no primary term	\$0.78m	\$1.47m	\$2.95m	26.4%
25% royalty, no primary term	\$1.06m	\$1.70m	\$0.71m	40.0%
39% royalty, no primary term	\$1.14m	\$1.60m	\$0.18m	50.8%
39% royalty, 5 yr. primary term	\$1.15m	\$1.62m	\$0.10m	49.6%
39% royalty, 3 yr. primary term	\$1.15m	\$1.63m	\$0.07m	49.3%

Note: "Baseline model" is our estimated model from Section 5. For cases with a primary term, upon expiration the original firm has the option to extend the lease by paying another bonus. Each extension involves the same royalty rate and term length as the original lease. Values are post-tax for a unit with mean OGIP, taken as expectations over $\psi(\theta)$ and all possible price paths starting from conditions in Q1 2009.

distortion to the drilling timing of the lease's first well. This effect also benefits the owner. Its value is greatest in the model with one well per lease and no endogenous inputs, since in that model the primary term directly addresses the royalty's only distortion, which is itself large due to the size of the royalty. Second, the primary term effectively increases dispersion in the type space at the time the lease is offered, since the deadline affects low θ types more adversely than high θ types. This effect is the opposite of that of the royalty and reduces the owner's value. Moreover, the owner responds to this effect by setting the bonus so that fewer types agree to the lease offer, reducing total surplus. This second effect outweighs the first in our baseline model but not when we set $\beta_w = 0$ and model only one well per unit.

We formalize this intuition using an alternative version of our model that is analytically tractable but abstracts away from some features of our computational model: it omits rig dayrates D_t and the cost shocks ν_{it}, and it assumes that the lease can only accommodate a single well. This model, which is presented in Appendix E, draws from Laffont and Tirole (1986) and Board (2007) to characterize a TIOLI menu of contracts that the owner can offer the firm in order to maximize its expected revenue. We find that the optimal contract includes a contingent payment paid at the time of drilling that consists of two components: a royalty on total revenue and a fixed payment that could be positive or negative. We show that if β_w is sufficiently small, the fixed payment is made from the owner to the firm and functions similar to a primary term: it incentivizes earlier drilling to counteract the delay effects of the high royalty rate. If β_w is sufficiently large, however, the fixed payment is made from the firm to the owner—in addition to a relatively small royalty—further delaying drilling and playing a role opposite that of a primary term.

8 Conclusion

This paper begins by presenting evidence that primary terms embedded into mineral leases in the Haynesville Shale in Louisiana have led to substantial bunching in the timing of firms' drilling activity. While this bunching is distortionary, it also hastens drilling, counteracting a delay distortion induced by lease royalties. We study these effects using an estimated model that incorporates two important features of the Haynesville Shale: the large impact of water input choices on production, and the fact that drilling one well on a unit gives the firm the option to drill follow-up wells. We find that a primary term can modestly increase the total expected surplus from a mineral lease with a 25% royalty rate. When we use an alternative model that shuts down these two features, we find that primary terms are considerably more beneficial. Finally, we find that when we model the owner's decision of what lease terms to offer, the effect of a primary term on the owner's expected revenue is quite small, and positive only when we use the alternative model.

These findings are related to recent litigation and institutional developments in the shale industry. In the Louisiana Haynesville, regulators impose 640-acre pooling unit sizes and have prevailed in court over mineral owners who argued that each well drilled should only hold the acreage that it drained rather than the entire unit (Gatti vs State of Louisiana 2014). In other states, however, large mineral owners have moved to include "retained acreage" clauses in their leases that strictly limit the acreage that any one well can hold.

This paper's model could be enriched in future work to more fully understand the economics of primary terms. Our model abstracts away from factors such as rig availability, learning, well interference, and financial frictions that might impact firms' drilling timing decisions (Hodgson 2018; Steck 2018; Agerton 2020; Gilje et al. 2020); integrating these features could reveal new implications for lease design. The contracting model under information asymmetry in Section 7 could be enriched with more balanced bargaining power and the possibility that the owner could contract with a different firm after the original primary term expires. Future work could also incorporate pre-specified lease extension options, which operate like one-time rental payments, encouraging earlier drilling and avoiding state-contingent renegotiation at the expiration of the primary term. Finally, work is needed to understand why lease terms appear to be "sticky", exhibiting little variation across space and time.

Extending our modeling framework to other settings is also likely to be worthwhile. Expanding the scope to other major shale plays would permit an assessment of how royalties and primary terms have affected aggregate U.S. oil and gas supply. Such work could examine the role of mineral leases in driving misallocation of shale drilling in the U.S., relating to Asker et al.'s (2019) work on aggregate wedges between optimal and observed global oil extraction and to Gilje et al.'s (2020) documentation of how debt renegotiations have distorted U.S. shale drilling. Our framework could also be used to evaluate the economics of carbon policies in a second-best environment in which oil and gas production is already distorted by mineral lease terms. Finally, the ideas in this paper could be extended to settings such as retail franchising or intellectual property licensing in which principals sell time-limited development options to agents.

References

Ackerberg, Daniel A. and Paul J. Devereux, "Improved JIVE Estimators for Overidentified Linear Models with and without Heteroskedasticity," *The Review of Economics and Statistics*, 2009, pp. 351–362.

Agerton, Mark, "Learning Where to Drill: Drilling Decisions and Geological Quality in the Haynesville Shale," 2020. working paper.

Anderson, Soren T., Ryan Kellogg, and Stephen W. Salant, "Hotelling Under Pressure," *Journal of Political Economy*, June 2018, *126*, 984–1026.

Asker, John, Allan Collard-Wexler, and Jan De Loecker, "(Mis)Allocation, Market Power, and Global Oil Extraction," *American Economic Review*, 2019, *109* (4), 1568–1615.

Bhattacharya, Vivek, Andrey Ordin, and James W. Roberts, "Bidding and Drilling Under Uncertainty: An Empirical Analysis of Contingent Payment Auctions," *Journal of Political Economy*, 2018, *130* (5).

Bloomberg, "U.S. Crude oil and natural gas futures prices," accessed 3 October, 2017.

Board, Simon, "Selling Options," *Journal of Economic Theory*, 2007, *136*, 324–340.

Brown, Jason, Timothy Fitzgerald, and Jeremy G. Weber, "Capturing Rents from Natural Resource Abundance: Private Royalties from U.S. Onshore Oil & Gas Production," *Resource and Energy Economics*, 2016, *46*, 23–38.

Bureau of Labor Statistics, "Consumer Price Index, all urban, all items less energy, not seasonally adjusted (CUUR0000SA0LE)," https://www.bls.gov/cpi/data.htm; accessed 8 July, 2018.

Cadotte, R. J., Andrew Whitsett, Melfard Sorrell, and Ben Hunter, "Modern Completion Optimization in the Haynesville Shale Modern Completion Optimization in the Haynesville Shale," in "SPE Annual Technical Conference and Exhibition" Society of Petroleum Engineers October 2017.

Covert, Thomas R., "Experiential and Social Learning in Firms: the Case of Hydraulic Fracturing in the Bakken Shale," 2015. working paper.

_ **and Richard L. Sweeney**, "Relinquishing Riches: Auctions vs Informal Negotiations in Texas Oil and Gas Leasing," 2019. working paper.

Enverus, "DI Classic completion, lease, and well data," accessed in batches on 1 June, 2012, 6 May, 2014, 7 January, 2016, 21 January 2016, and 22 February, 2016.

_ , "DI Desktop monthly well production data," accessed on 5 February, 2016.

_ , "Rig Day Rate Report," accessed on 2 November, 2017.

GAO, "Spectrum Management: FCC's Use and Enforcement of Buildout Requirements," 2014. GAO-14-236.

Gatti vs State of Louisiana, 2014. Louisiana Court of Appeals, CA 0289.

Gilje, Erik P., Elena Loutskina, and Daniel Murphy, "Drilling and Debt," *Journal of Finance*, 2020.

Gülen, Gürcan, Svetlana Ikonnikova, John Browning, Katie Smye, and Scott W. Tinker, "Production Scenarios for the Haynesville Shale Play," *SPE Economics & Management*, 2015, *7* (5).

Haile, Philip, Kenneth Hendricks, and Robert Porter, "Recent U.S. Offshore Oil and Gas Lease Bidding: A Progress Report," *International Journal of Industrial Organization*, 2010, *28*, 390–396.

Hendricks, Kenneth, Robert H. Porter, and Guofo Tan, "Optimal Selling Strategies for Oil and Gas Leases with an Informed Buyer," *American Economic Review Papers and Proceedings*, May 1993, *83* (2), 234–239.

Hiller, Jennifer, "Eagle Ford Lease Deadlines Driving Drilling," *San Antonio Express-News*, 25 Nov. 2012.

Hodgson, Charles, "Information Externalities, Free Riding, and Optimal Exploration in the UK Oil Industry," 2018. working paper.

Kaiser, Mark J., "Profitability Assessment of Haynesville Shale Gas Wells," *Energy*, 2012, *38* (1), 315–330.

Kalnins, Arturs, "Overestimation and Venture Survival: an Empirical Analysis of Development Commitments in International Master Franchising Ventures," *Journal of Economics and Management Strategy*, Winter 2005, *14* (4), 933–953.

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

_ **and Mar Reguant**, "Energy and Environmental Markets, Industrial Organization, and Regulation," in Ali Hortaçsu, Kate Ho, and Alessandro Lizzeri, eds., *Handbook of Industrial Organization*, Vol. 5, Elsevier, 2021, pp. 615–742.

Kolesàr, Michal, "Estimation in an Instrumental Variables Model with Treatment Effect Heterogeneity," 2013. Working paper.

Kong, Yunmi, Isabelle Perrigne, and Quang Vuong, "Multidimensional Analysis of Contracts: An Empirical Analysis," *American Economic Review*, 2022, *112* (5), 1703–1736.

Laffont, Jean-Jacques and Jean Tirole, "Using Cost Observation to Regulate Firms," *Journal of Political Economy*, Jun. 1986, *94* (3), 614–641.

Li, Qi and Thanasis Stengos, "Semiparametric Estimation of Partially Linear Panel Data Models," *Journal of Econometrics*, 1996, *71*, 389–397.

Litwak, Mark, *Contracts for the Film & Television Industry, 3rd Edition*, Silman-James Press, 2012.

Louisiana Department of Natural Resources, "Well and unit shapefiles," http://sonris-www.dnr.state.la.us/gis/agsweb/IE/JSViewer/index.html?TemplateID=181; accessed 7 January, 2016.

_ , "Well drilling and completion cost reports," http://sonris.com/; accessed in spring and summer, 2016.

_ , "Well tabular data," http://sonlite.dnr.state.la.us/pls/apex/f?p=108:2:32690097475501::::::; accessed in May, 2016.

Metcalf, Gilbert E., "Investment in Energy Infrastructure and the Tax Code," in Jeffrey R. Brown, ed., *Tax Policy and the Economy, Volume 24*, University of Chicago Press, 2010.

Montgomery, Carl, "Fracturing Fluids," in Andrew P. Bunger, John McLennan, and Rob Jeffrey, eds., *Effective and Sustainable Hydraulic Fracturing*, IntechOpen, 2013.

Newell, Richard G., Brian C. Prest, and Ashley Vissing, "Trophy Hunting vs. Manufacturing Energy: The Price-Responsiveness of Shale Gas," *Journal of the Association of Environmental and Resource Economists*, 2019, *6*, 177–217.

Olley, G. Steven and Ariel Pakes, "The Dynamics of Productivity in the Telecommunications Equipment Industry," *Econometrica*, 1996, *64* (6), 1263–1297.

Ordin, Andrey, "Investment and Taxation: the Case of Oil and Gas in the Permian Basin," 2019. working paper.

Patzek, Tad W., Frank Male, and Michael Marder, "Gas Production in the Barnett Shale Obeys a Simple Scaling Theory," *PNAS*, Dec. 2013, *110* (49), 19731–19736.

Penny, Glenn S., Michael W. Conway, and Wellington Lee, "Control and Modeling of Fluid Leakoff During Hydraulic Fracturing," *Journal of Petroleum Technology*, June 1985, *37* (6), 1071–1081.

Robinson, Peter M., "Root-n-consistent Semiparametric Regression," *Econometrica*, Jul. 1988, *56* (4), 931–954.

Rust, John, "Optimal Replacement of GMC Bus Engines; An Empirical Model of Harold Zurcher," *Econometrica*, 1987, *55* (5), 999–1033.

Skrzypacz, Andrzej, "Auctions with Contingent Payments—an Overview," *International Journal of Industrial Organization*, 2013, *31*, 666–675.

Steck, Andrew L., "Industry Dynamics with Social Learning: Evidence from Hydraulic Fracturing," 2018. working paper.

Ward, Melissa M. and Timothy W. Kelsey, "Natural Gas Lessors' Experiences in Bradford and Tioga Counties, 2010," 2011. Marcellus Education Fact Sheet, PennState Extension.

Yarushina, Viktoriya M., David Bercovici, and Michael L. Oristaglio, "Rock Deformation Models and Fluid Leak-Off in Hydraulic Fracturing," *Geophysical Journal International*, 2013, *194*, 1514–1526.

Online Appendix for "Drilling Deadlines and Oil and Gas Development"

A Detail on data sources, cleaning, and merging

In this data appendix, we discuss: (1) our data sources; (2) how we estimate well decline and the present value of well cumulative production; (3) how we clean lease data and match leases to units; and (4) how we match wells to Haynesville units.

A.1 Data sources

We gather data from the following sources:

- Publicly-available Louisiana DNR Strategic Online Natural Resources Information System (SONRIS) data on well drilling and completions.
  - Shapefiles from Louisiana Department of Natural Resources (2016a) that include wells' top hole location, bottom hole location, and lateral location; and units' boundaries.
  - Well tabular data from Louisiana Department of Natural Resources (2016c) that include spud date, completion date, well name, and formation targeted.
  - Drilling cost and fracking input data from Louisiana Department of Natural Resources (2016b). 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. We used manual double-entry to digitize this information from the raw pdf files.
- Publicly-available Louisiana parish boundaries from US Census Bureau (2020).
- Enverus well and completion shapefiles, from Enverus (2016a).
- Enverus production data from Enverus (2016b). Enverus takes unit-level reported monthly production data from the Louisiana DNR and then imputes well-level monthly production using the start date of each well's production.
- Enverus lease data, from Enverus (2016a). Enverus collects data on leases signed in Louisiana. Further details are below.
- Enverus dayrate data from Anderson et al. (2018) and Enverus (2017). 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, from Bloomberg (2017). We deflate these prices, rig dayrates, and drilling cost data to December 2014 dollars using the Bureau of Labor Statistics'

Consumer Price Index for all goods less energy, all urban consumers, and not seasonally adjusted (Bureau of Labor Statistics 2018). The CPI series ID is CUUR0000SA0LE.

- Original gas in place (OGIP) from Gülen et al. (2015).

A.2 Well data and production decline estimation

We identify Haynesville wells using three sources. First, for each well we see if the well is included in an auxiliary DNR file that is limited to Haynesville wells. Second, we use the name of the well: wells targeting the Haynesville typically have names that begin with "HA" or "HAY". Third, we check whether the listed formation that the well targets is the Haynesville. We denote a well as a Haynesville well if it satisfies at least one of these three criteria. We also impose a restriction that Haynesville wells must have been spudded on or after September 2006.

To estimate wells' production decline, we follow Patzek et al. (2013), which derives decline curves for shale gas formations. This paper shows that production initially declines inversely proportional to the square root of time, and then begins to decline more quickly, at an exponential rate, once the well's fractures interfere with one another. More precisely, we assume that cumulative production of natural gas for well j at month t takes the following functional form:

$$m_j(t) = \begin{cases} M_j\sqrt{t/\tau} & \text{if } 0 \leq t \leq \tau \\ M_j + \dfrac{M_j}{2\tau d}\left[1 - \exp(-d(t-\tau))\right] & \text{if } t > \tau \end{cases} \tag{A.1}$$

where τ is the time at which the decline function changes to exponential, d is the exponential decline rate, and M_j is a well-specific production multiplier corresponding to the expected cumulative production at $t = \tau$.

Before estimating these parameters, we make a number of adjustments to the data. First, because a well's production is substantially affected by the length of the lateral well leg, we normalize the measure of cumulative production by a scalar s_j 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.

We find that about 7% of wells had recompletions. As recompletions are designed to rapidly increase production, we exclude from the data observations that come during or after months in which well recompletions were performed. (When we later use our estimates to predict total production, we assume no recompletions.)

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 noisy. This noise is due in part to the fact that hydrofracturing water is still being back-produced in the early months of production. Similarly, for any given date with no production but in which 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, which finds 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 (A.2).

$$\sum_j \sum_{t|t\geq 4} (\log m_j(t)s_j - \log \widehat{m}_j(t|M_j, \tau, d))^2 \tag{A.2}$$

The estimated decline parameter d is equal to 0.037. The 25th, 50th, and 75th percentiles of the estimated M_i are 1.57 million, 2.09 million, and 2.63 million mmBtu, respectively.

We then use our estimates of d and M_j to predict total discounted well production (Equation (A.3)). Following Gülen et al. (2015), we assume that wells have a total production lifetime of 20 years. We use an annual discount factor of 0.909, following Kellogg (2014). In Panel (a) of Figure 1, we map our measures of the present value of total well production. Where there are multiple wells, we take an average over all wells within the unit. Units with no drilling have no shading and are labeled NA.

$$Y_j = \sum_{t=1}^{240} [\widehat{m}_j(t|M_j, \tau, d) - \widehat{m}_j(t-1|M_i, \tau, d)]\delta^{t-1} \tag{A.3}$$

A.3 Lease data and clustering of duplicate leases within units

In February 2016, we downloaded raw data of oil and gas leases in Louisiana from Enverus. 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. Because we ultimately map leases to units, 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. 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 from our sample.

We drop leases with zero or missing acreage. We also drop excess lease observations that are perfect duplicates, leases that have lengths of fewer than 10 days, and leases in which the reported township, range, and section are not within the stated reported parish.

We 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 these leases by identifying duplicates that share the same grantee name and the same acreage, and where the acreage reported is unusual—i.e., is either large and/or is not equal to a multiple of common lot sizes (e.g., 10 acres, 40 acres). In these cases, we impute a new acreage measure by dividing the reported acreage by the number of apparent duplicates. In cases where a lease spans multiple sections and acreage within each section is not reported, we assume total acreage is divided equally between the spanned sections.

After taking the above steps, we find that the total leased acreage in a unit still sometimes adds up to more than the total acreage of the unit (usually 640 acres), and sometimes significantly so. Many of these cases appear to be driven by undivided mineral interests: cases where there are multiple grantors on the same plot (e.g., husband and wife, multiple siblings, or cousins), and separate observations for each grantor. In other cases, it appears that data were entered multiple times and inconsistencies were not reconciled, so that the excess observations were not dropped when we removed duplicates.

To identify these likely duplicates, we use an agglomerative, 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 obtain 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 in which the distance between observation i and observation j takes the form:

$$d_{ij} = \sqrt{\sum_k w_k m_k(x_i^k, x_j^k)} \tag{A.4}$$

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 technique is useful for catching cases with identical last names but differing or missing 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. We set $w_k = 1$ for all characteristics other than acreage, for which we set $w_k = 100$. This weighting ensures that leases that vary in acreage will not be presumed to be duplicates.

How many observations are clustered together depends on the threshold level of similarity imposed by the researcher. To determine our threshold, we choose a calibration date of January 1, 2010, examining only the leases that were active on that date.[1] We first examine every possible threshold that could be used to cluster the leases in each unit. For each possible threshold, we find the resulting clusters, downweight each lease's acreage by the total number of observations in its cluster, and then compute what total leased acreage would be within the unit. Then, for each unit, we find the threshold would be that would set total acreage leased to be equal to or just less than the total unit area. We refer to this threshold as the unit-level threshold height. We then set our preferred overall threshold to be equal to the 90th percentile of all the unit-level thresholds. The threshold height that results from this computation is 1.644. We find similar results if we use a threshold height using the 85th percentile.

We then apply the clustering procedure, using this threshold, to each unit and each quarterly date of our sample, ranging from January 1, 2005, to January 1, 2016. This procedure 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 these situations, we take the inverse of the arithmetic average of the inverse downweight over all quarterly dates to obtain a master downweight for each lease (yielding, in this example, a weight of 1/5.5).

In some outlier unit-quarters we find that even with this downweighting, total leased acreage still exceeds section acreage. In these cases, we then proportionally reduce the area of all leases in the unit so that total leased acreage is equal to total unit acreage in the most heavily leased quarter.

A.4 Matching of wells to units

To match wells to units, we use information on the reported laterals, reported bottom holes, and reported top hole locations. If, for a given well, the data only report top hole location, we use the location of the top hole to identify which unit the well is in. If the data report bottom hole but not lateral information, we use the location of the bottom hole to identify which unit the well is in. If the data report lateral information, we use the unit that the lateral runs through to identify the well's unit.

In a few instances, the well lateral intersects multiple units. There are two possible reasons for these occurrences. One is that the well's top hole is located in a different unit than the unit the well extracts from, 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 unit. A

[1] We 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 in which 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 A.1: Estimates from bunching analysis

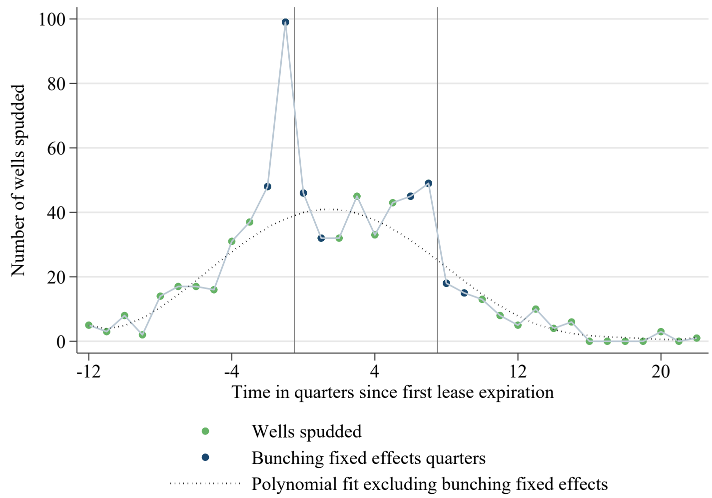

Note: Figure presents data and estimates corresponding to the bunching analysis in column (1) of Table A.I. Plotted data include counts of wells spudded, the quarters to which we add bunching fixed effects, and the polynomial predicted probabilities given the bunching estimator fixed effects. Timing 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.

second reason is that the well actually targets multiple units. In cases where a well bore passes through multiple units, we only match a well to a unit if at least 300 meters of the horizontal well bore pass through the unit.

B Additional empirical analysis

This appendix presents additional empirical results related to the bunching analysis presented in Section 4.

B.1 Bunching analysis

To test the statistical significance of the drilling bunching shown in Figure 5, we use a bunching estimator 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 (across all units in our analysis sample) for each quarter (34 quarters in total). 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 post_ext2).

Table A.I: Bunching estimates

	count	log(count)	count	log(count)	count	log(count)
pre_2	12.72	0.40	1.06	0.19	1.02	0.21
	(5.25)	(0.21)	(0.88)	(0.30)	(0.82)	(0.30)
pre_1	60.94	1.10	3.91	0.78	3.99	0.88
	(5.71)	(0.21)	(1.36)	(0.29)	(1.29)	(0.29)
post_1	6.03	0.31	0.31	-0.07	0.44	0.04
	(5.74)	(0.21)	(0.77)	(0.28)	(0.70)	(0.27)
post_2	-8.92	-0.09	-0.50	-0.21	-0.53	-0.17
	(5.32)	(0.19)	(0.61)	(0.25)	(0.58)	(0.25)
pre_ext2	13.70	0.26	1.24	0.52	1.39	0.53
	(4.75)	(0.19)	(1.11)	(0.34)	(1.05)	(0.34)
pre_ext1	21.73	0.52	1.79	1.19	1.93	1.15
	(5.23)	(0.22)	(1.26)	(0.31)	(1.16)	(0.31)
post_ext1	-4.99	-0.24	-0.32	-0.05	-0.23	0.04
	(5.19)	(0.21)	(0.72)	(0.32)	(0.72)	(0.32)
post_ext2	-3.70	-0.11	-0.26	0.28	-0.15	0.41
	(4.61)	(0.17)	(0.64)	(0.30)	(0.67)	(0.37)
Quarter of lease data	X	X				
Quarter by quarter of lease data			X	X	X	X
Calendar quarter fixed effects					X	X
R Squared	0.97	0.93	0.24	0.3	0.42	0.44
Observations	35	30	363	235	363	235

Note: Table presents estimates of Equation (A.5). Newey-West standard errors, computed with two quarterly lags, are in parentheses. Estimates in columns (1) and (2) use data that are aggregated to the lease-level quarter, which is defined as the time between first primary term expiration and spudding, measured at quarterly intervals. Estimates in columns (3) through (6) use data that are aggregated to the lease-level quarter by calendar-level quarter. That is, these columns aggregate wells drilled that share both a common lease-level quarter and a common calendar quarter-of-sample. Estimates in columns (5) and (6) include calendar quarter fixed effects.

We then estimate a regression of the form:

$$\begin{aligned} c_t = f(t) &+ \beta_1 \cdot \text{pre\_2} + \beta_2 \cdot \text{pre\_1} + \beta_3 \cdot \text{post\_1} + \beta_4 \cdot \text{post\_2} \\ &+ \beta_5 \cdot \text{pre\_ext2} + \beta_6 \cdot \text{pre\_ext1} + \beta_7 \cdot \text{post\_ext1} + \beta_8 \cdot \text{post\_ext2} + \varepsilon_t \end{aligned} \tag{A.5}$$

where c_t is total well count, t is quarter, and $f(t)$ is a polynomial of degree 9. Our main regression estimates are in column (1) of Table A.I. The estimates of β_1, β_2, β_5, and β_6 are all statistically significant with p-values less than 0.05, indicating 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. In Column (2), we present results from the same empirical specification as column (1) except that the dependent variable is the log of the count rather than the count, and results are similar. For a sense of magnitude, the estimate of β_2 (the coefficient on "pre_1") in column (2) means that the actual number of wells drilled is 1.1 log points larger than the polynomial fit in the quarter prior to expiration. In Figure A.1, we plot our data and the number of wells predicted by our polynomial fit, which graphically displays the size of the bunching effect.

One might worry that periods with substantial lease expirations coincide with periods in which gas prices or industry-wide productivity is high. To address this possibility, we construct a measure of total spud counts at the calendar quarter by quarter of lease level. For example, one observation in this count data will be the total number of spuds in 2010 quarter 3 when the spud happened between 3 and 6 months before the first primary term is set to expire. In columns (3) and (4) of Table A.I, we present estimates from the same empirical specifications as columns (1) and (2), only with this more disaggregated data. Columns (5) and (6) then add in calendar-time quarter fixed effects. Across columns (3) and (5), the coefficient on pre_1 is large, statistically significant, and similar in magnitude, implying that the high drilling before the expiration date is not being driven by high drilling at particular calendar dates. The same holds in logs for columns (4) and (6).

B.2 Additional descriptive figures and table

Figure A.2: Comparison of units with identical vs different neighboring operators

(a) Units in analysis sample

(b) All Haynesville units

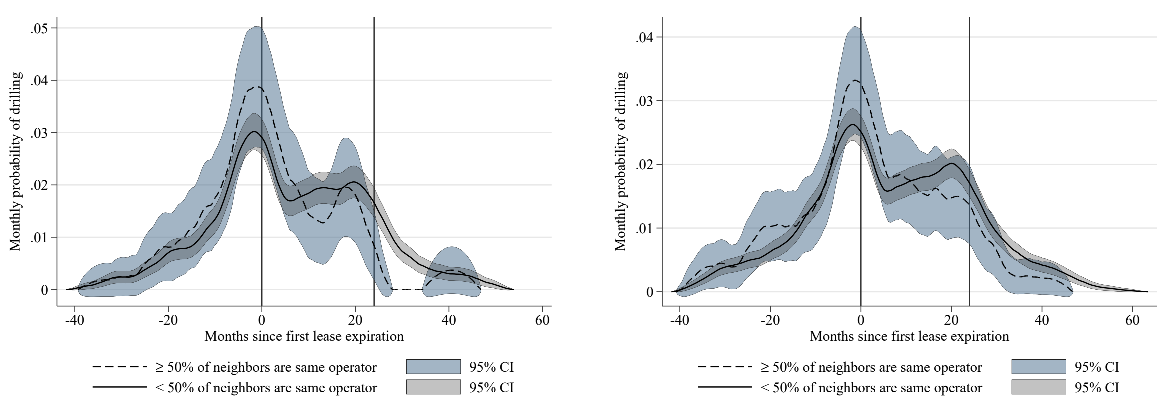

Note: Figure shows kernel-smoothed estimates of the probability of drilling the first Haynesville well in a unit on a given date, relative to the expiration date of the first lease within the unit to expire. Vertical lines are drawn at the date of first lease expiration and two years after first lease expiration. Figure compares units in which $\geq 50\%$ of the nearby units have the same operator vs. units where $\geq 50\%$ 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 (results are similar if we use a threshold of 1.7 miles, which will include the diagonal units). In Panel (a), we limit the units to our analysis sample, as described in subsection 3.4. In Panel (b), we show results using all Haynesville units.

Figure A.3: Wells' production, water use, and cost vs. time relative to first lease expiration

(a) First 12 months of gas production (log mmBtu)

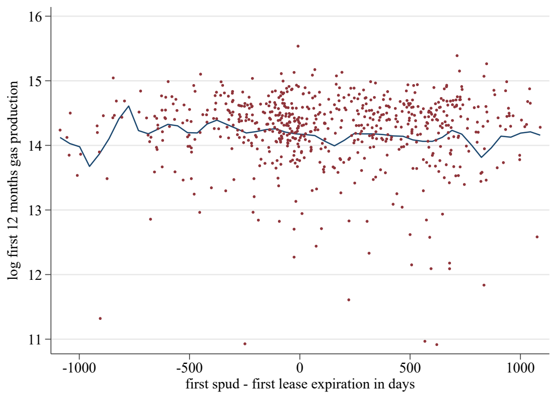

(b) Water use (log gallons)

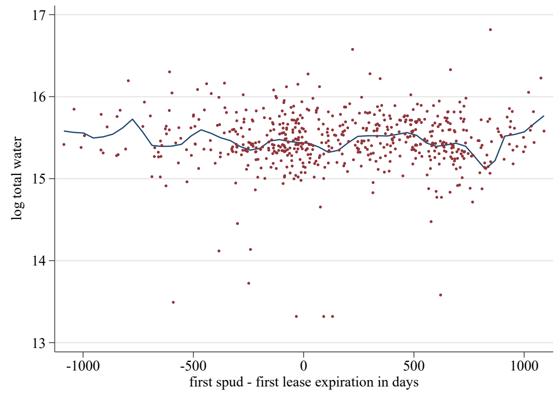

(c) Reported drilling and completion cost (log $2014)

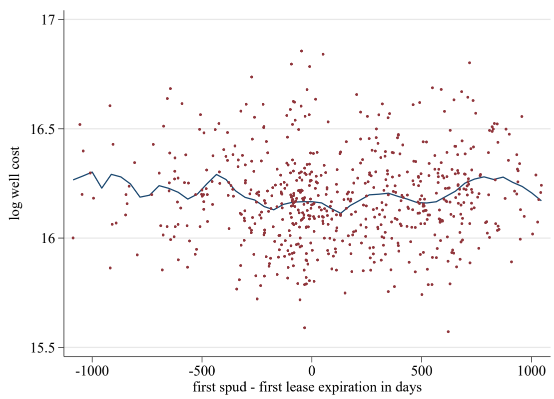

Note: Figure plots natural gas production, water input, and reported drilling cost for the first well drilled in each unit against the well's spud date relative to the date of first lease expiration (measured in days). The line is the predicted value from a local polynomial regression.

Figure A.4: Log water use versus time

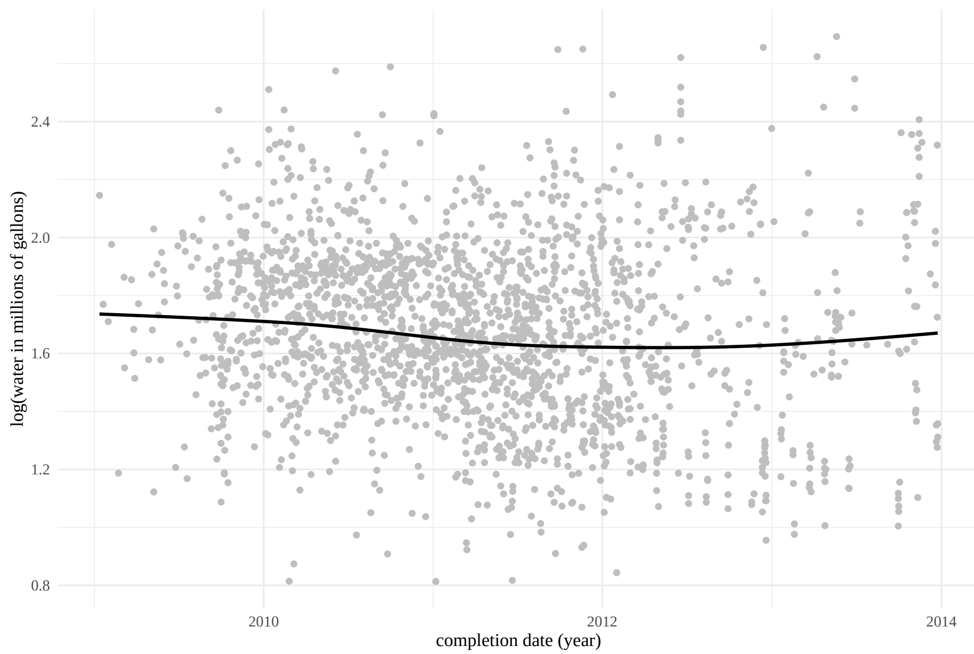

Note: Figure presents a scatter plot and lowess estimate of each well's log water use versus its spud date. Scatter plot excludes a small number of outlier water observations.

Table A.II: Regressions of lease terms on OGIP and natural gas prices

	(1) Royalty	(2) Term	(3) 1(Extension)
OGIP	0.0136	-0.0192	-0.0024
	(0.0040)	(0.0060)	(0.0004)
12 month NG futures (real)	-0.3832	0.1569	-0.0263
	(0.0667)	(0.0865)	(0.0065)
Constant	24.2351	36.1959	1.2124
	(0.7925)	(0.6830)	(0.0821)
R squared	0.0931	0.0168	0.0453
Observations	29370	37528	37448

Note: Regressions where the unit of observation is the lease and observations are weighted by lease acreage. Natural gas price is that at the time of lease signing. Standard errors are clustered by unit. Royalty is measured in percentages, e.g., a value of 25 for the dependent variable means a royalty of 25%. According to these estimates, moving from the lowest to highest OGIP unit in our data would imply an increase in the royalty of 2.4 percentage points, a decrease in the primary term of 3 months, and a 42 percentage point decrease in the probability of an extension.

C Details of model simulation and estimation

This appendix describes in greater detail how we simulate and estimate the model presented in Section 5.

C.1 Estimating β_w

We first discuss our bandwidth selection to nonparametrically control for latitude and longitude, and then discuss IJIVE, UJIVE, 2SLS, and OLS.

Bandwidth: Our IJIVE as well as UJIVE, 2SLS, and OLS estimates rely on non-parametric regression to control for latitude and longitude. To do this, we use a Gaussian kernel $\phi(0, \sigma_\phi^2)$ where ϕ is the normal pdf with a mean of 0 and a standard deviation of σ_ϕ. To calculate the bandwidth σ_ϕ, we use an optimal bandwidth approach using a leave-one-out estimator: We use the average log water for all wells j' other than well j to predict log water for well j. Our estimate of σ_ϕ is the value that minimizes the sum of the squares of the differences between the actual and predicted values of log water for well j:

$$\sigma_\phi = \arg\min_{\sigma>0} \sum_j \left(\text{log water}_j - \frac{\sum_{j' \neq j} \phi(d(j, j'), \sigma^2) \cdot \text{log water}_{j'}}{\sum_{j' \neq j} \phi(d(j, j'), \sigma^2)} \right)^2 \tag{A.6}$$

In our regressions (including residualization for IJIVE as well as Robinson-semiparametric 2SLS and OLS regressions—but not during cross-validation), we also apply a caliper that assigns zero weight to any observations greater than four bandwidths away from the observation of interest.

We calculate σ_ϕ = 3,632 meters. If we instead use production rather than log water as an input to calculate σ_ϕ, we calculate σ_ϕ = 2,407 meters. Our IJIVE and UJIVE estimates of β_w are quantitatively similar using this smaller bandwidth.

IJIVE: Following Ackerberg and Devereux (2009), we identify β_w using the following steps:

1. For each variable—production, log water, and each of the month effects—we project the variable on a non-parametric function of latitude and longitude using the Gaussian kernel $\phi(0, \sigma_\phi^2)$. We then compute the residualized variable as the difference between the variable and the predicted value of the variable.
2. For each well j, we use OLS to project residualized log water for well j on residualized month fixed effects for all wells other than well j. We use this projection to compute predicted residualized log water for well j.
3. We calculate β_w by projecting residualized production on predicted residualized log water using OLS.

UJIVE: Following Kolesàr (2013), the UJIVE estimation of β_w takes a similar approach to IJIVE in that it uses leave-one-out prediction and partials out the contribution of geology. However, it does so by constructing a new instrument Z_j that is then used in conventional 2SLS:

1. For each well j, use all wells other than well j and project log water on a non-parametric function of latitude and longitude as well as month fixed effects. Compute predicted log water for well j: $\widehat{\log w_j}$.

2. For each well j, use all wells other than well j and project log water on only a non-parametric function of latitude and longitude. Computed predicted log water for well j: $\widetilde{\log w_j}$.
3. Construct the UJIVE instrument as $Z_j = \widehat{\log w_j} - \widetilde{\log w_j}$.
4. First stage: Use Robinson (1988) to project log water on Z_j and a non-parametric function of latitude and longitude. Use this projection to construct predicted log water.
5. Second stage: Use Robinson (1988) to project production on predicted log water and a non-parametric function of latitude and longitude. The coefficient on predicted log water is β_w.

2SLS: Our 2SLS specification uses month fixed effects as instruments, in two steps:

1. Use Robinson (1988) to project log water on month fixed effects and a non-parametric function of latitude and longitude. Use this projection to construct predicted log water.
2. Use Robinson (1988) to project production on predicted log water and a non-parametric function of latitude and longitude. The coefficient on predicted log water is β_w.

OLS: Our OLS specification is the Robinson (1988) double-residual regression projecting production on log water and flexibly controlling for latitude and longitude.

C.2 Profits and optimal water input

The static drilling profits accruing to the firm differ depending on whether its profits before taxes, royalties, and operating costs are positive (i.e., whether the well "pays out"). As we discuss in Section 2, unleased mineral interests are not liable for well costs if the well fails to pay out. In addition, severance taxes are waived. These rules create a kink in the profit function at the payout point. If the well pays out, profits are given in Equation (2) in the main text. If the well does not pay out, profits are given by Equation (A.7):

$$\pi^{-}_{ijt} = (1-\tau)(1-f_{it}k_i - c)P_t g(\theta_i, X_i, W_j, \varepsilon_{it}) - (1-\tau)P_w W_j + (1-\tau_c)(\alpha_0 + \alpha_1 D_t) \quad \text{(A.7)}$$

This kink also affects the optimal amount of water use, since the firm's first-order condition that determines optimal water use will depend on whether the firm expects the well to pay out or not. Optimal water use W^*_{it} is given by Equation (A.8) if the well pays out:

$$\log(W^*_{ijt}) = \log(\beta_w) + \log P_t - \log(P_{wt}) + \log\left(\frac{(1-s)(1-k_i) - c}{1-s+sk_i}\right) \quad \text{(A.8)}$$

and by Equation (A.9) if it does not pay out:

$$\log(W^*_{ijt}) = \log(\beta_w) + \log P_t - \log(P_{wt}) + \log(1 - f_{it}k_i - c). \quad \text{(A.9)}$$

Given the parameter inputs to the static profit function, we determine the optimal water use W^*_{it} by first finding the values W_+ and W_- that solve equations (A.8) and (A.9), respectively. If W_+ results in positive payout, we set $W^*_{ijt} = W_+$. Alternatively, if W_- results in negative payout, we set $W^*_{ijt} = W_-$. Finally, it is possible that W_- results in positive payout while W_+ results in negative payout. In that case we interpolate the value of $W^* \in (W_+, W_-)$ that results in zero payout.

Table A.III: OLS estimates of Equation (A.10)

log NG price $(1-\gamma_1)$	-0.006
	(0.185)
log dayrate $(-\gamma_2)$	-0.24
	(0.197)
intercept $(log(\beta_w)-\gamma_0)$	4.472
	(4.449)
R^2	0.006
N	2,019

Note: The dependent variable is $\log(W_i)-\log\left(\frac{(1-s)(1-k_i)-c}{1-s+sk_i}\right)$. Sample uses all wells in the production estimation sample. Standard errors are clustered at the township level.

C.3 Water price estimation

Once the production function coefficient β_w is estimated, each term in the first-order condition for optimal water use is known except for the price of water P_{wt}. We use this fact to estimate the γ parameters in the water price projection (Equation (5) in the main text) by combining Equation (5) and Equation (A.8) into Equation (A.10), which we estimate by OLS:[2]

$$\log(W_j)-\log\left(\frac{(1-s)(1-k_i)-c}{1-\tau+sk_i}\right)=(\log\beta_w-\gamma_0)+(1-\gamma_1)\log P_t-\gamma_2\log D_t+\omega_j. \quad (A.10)$$

Estimates from Equation (A.10) are presented in Table A.III. Estimates of γ_1 and γ_2 are similar if we include nonparametric controls for latitude and longitude rather than an intercept term. Because our estimate of β_w is used to identify γ_0, and because we use the same sample to estimate β_w as to estimate the water price process, we use bootstrapping with 5,000 draws to estimate the variance-covariance matrix of $[\beta_w,\gamma_0,\gamma_1,\gamma_2]$. Our bootstrap draws are clustered at the township level to account for spatial correlation.

To back out the time series of water price shocks ω_t (defined in Equation (5) in the main text), we use the fact that the expected within-month mean of the residuals ω_j from Equation (A.10) are the negative of the ω_t. Because actual water use is noisy, we obtain our estimates of ω_t by applying Bayesian shrinkage to the within-month mean of the residuals ω_j.

We graph the resulting estimated implied log water price series in Figure A.5, aggregated to the quarterly level. As implied by our large and positive estimate of γ_1, water prices follow natural gas prices closely. These estimated prices should be thought of as the marginal cost of not just the water itself but also the associated labor and capital (e.g., pumping equipment) necessary to conduct the fracturing job.

[2]We use the first order condition for when the well "pays out" (as opposed to Equation (A.9)).

Figure A.5: Estimated log water price relative to natural gas price, dayrate, and average log water use

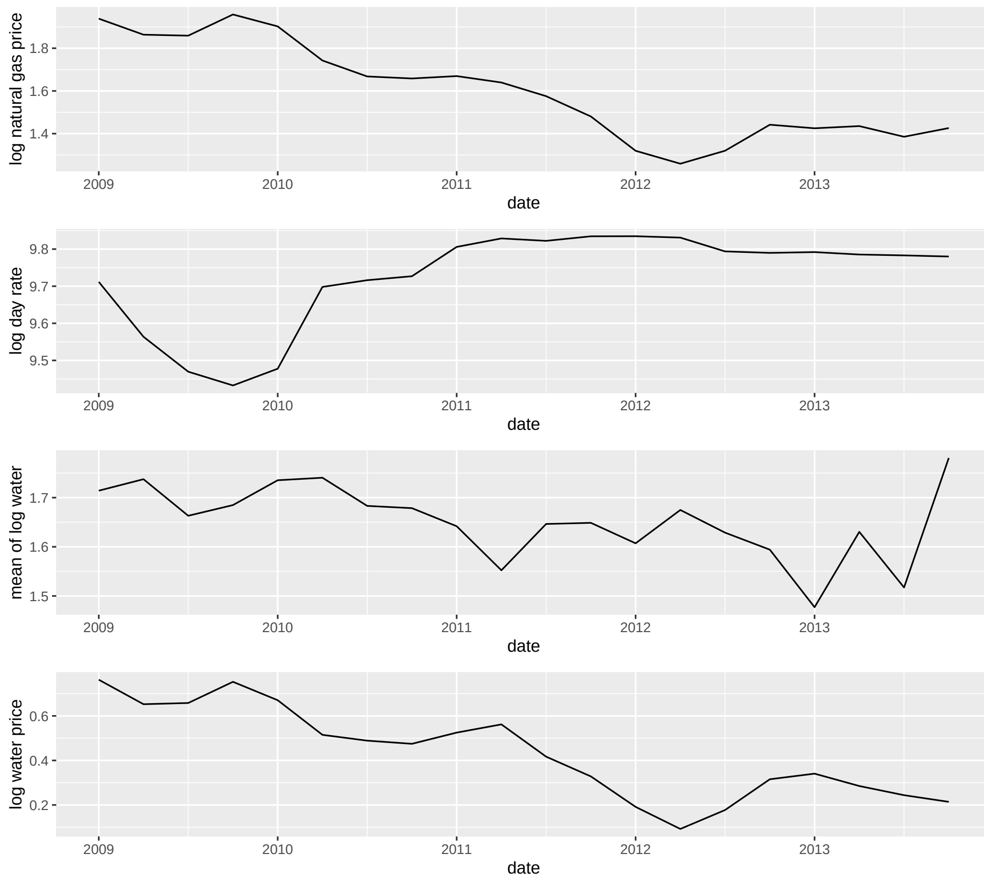

Note: The natural gas price is in \$/mmBtu, the dayrate is in \$/day, water use is in millions of gallons, and the water price is in \$/gallon.

C.4 Value function and model of predicted drilling probabilities

Our assumption in Equation (6) that expected profits are additive in the number of wells and in the cost shocks ν_{it}^1 and ν_{it}^0 allows us to significantly decrease the number of choices we need to consider. Flow profits increase linearly in the number of wells drilled, so that if the firm faces an infinite horizon problem, the firm will always drill either zero wells or all possible M wells. If the unit's leases have primary terms and are not held by production, however, the firm may prefer to drill a single well that extends any remaining leases indefinitely. Beyond that first well, the incremental payoffs are again constant, so the firm would never choose to drill strictly between one and M wells.

We formalize this concept and derive the probabilities of drilling 0, 1 and M wells below. 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.

- Let $E[\tilde{V}_{i,t+1}|S_{it},\theta_i] \equiv (1/M)E[V_{i,t+1}|S_{it},\theta_i]$ denote the firm's per-well continuation value at t if the lease has not been held by production at period t.
- Let $E[\tilde{U}_{i,t+1,t}|S_{it},\theta_i]$ denote the firm's per-well continuation value if the lease has been held by production on date t (i.e., if first drilling activity took place at time t).

Because the constraint of a primary term decreases the expected continuation value of the unit, $E[\tilde{V}_{i,t+1}|S_{it},\theta_i] < E[\tilde{U}_{i,t+1,t}|S_{it},\theta_i]$.

We focus on the firm's choice set for time periods when drilling has not occurred (as our maximum likelihood estimation only uses information on the timing of the first well), putting aside for the moment the cost shocks ν_{it}. The firm has three choices. First, the firm may choose to drill all M wells ($m = M$), thus ending the optimal stopping problem. The firm receives M times the per-well profits, where the expectation denotes the expectation over water prices P_{wt}:

$$V_{i,t}^M(S_{it},\theta_i) = ME_{P_{wt}}[\pi_{it}(S_{it},\theta_i)]. \tag{A.11}$$

Second, the firm may choose to drill only one well ($m = 1$), thus holding the unit. The firm then receives the profits from that well and the continuation payoff from the option to drill $M-1$ future wells:

$$V_{i,t}^1(S_{it},\theta_i) = E_{P_{wt}}[\pi_{it}(S_{it},\theta_i)] + \delta(M-1)E[\tilde{U}_{i,t+1,t}|S_{it},\theta_i] \tag{A.12}$$

Finally, the firm can retain the lease without drilling ($m = 0$) and receive the continuation value associated with an undrilled unit capable of holding M wells.

$$V_{i,t}^0(S_{it},\theta_i) = \delta ME[\tilde{V}_{i,t+1}|S_{it},\theta_i] \tag{A.13}$$

Now we incorporate the additive cost shocks. For each well that the firm drills, it gets the shock ν_{it}^1, and for each potential well that it does not drill, it gets ν_{it}^0 (scaled by the leased acreage share f_{it}). Thus, the full static payoffs from the three choices are:

$$\begin{aligned} \text{Drill all } M \text{ wells: } & V_{i,t}^M(S_{it},\theta_i) + f_{it}M\nu_{it}^1 \\ \text{Drill 1 well: } & V_{i,t}^1(S_{it},\theta_i) + f_{it}[\nu_{it}^1 + (M-1)\nu_{it}^0] \\ \text{Drill zero wells (continue): } & V_{i,t}^0(S_{it},\theta_i) + f_{it}M\nu_{it}^0 \end{aligned} \tag{A.14}$$

This structure leads to a tractable set of choice probabilities. Combining equations (A.14) with equations (A.11)–(A.13), we can characterize the firm's choice probabilities. The firm prefers drilling all M wells to a single well if:

$$\nu_{it}^1 - \nu_{it}^0 > \frac{1}{f_{it}}\left[\delta E[\tilde{U}_{i,t+1,t}|S_{it},\theta_i] - E_{P_{wt}}[\pi_{it}(S_{it},\theta_i)]\right]. \tag{A.15}$$

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

$$\nu_{it}^1 - \nu_{it}^0 > \frac{1}{f_{it}}\left[\delta E[\tilde{V}_{i,t+1}|S_{it},\theta_i] - E_{P_{wt}}[\pi_{it}(S_{it},\theta_i)]\right]. \tag{A.16}$$

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

$$\nu_{it}^1 - \nu_{it}^0 > \frac{1}{f_{it}} \left[M\delta E[\tilde{V}_{i,t+1}|S_{it},\theta_i] - (M-1)\delta E[\tilde{U}_{i,t+1,t,\theta}|S_{it},\theta_i] - E_{P_{wt}}[\pi_{it}(S_{it},\theta_i)] \right] \quad \text{(A.17)}$$

Since $E[\tilde{U}_{i,t+1,t}|S_{it},\theta_i] > E[\tilde{V}_{i,t+1}|S_{it},\theta_i]$, 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. By the same inequality, if a firm prefers drilling zero wells to drilling one well, it also prefers to drill zero wells over drilling M wells. Therefore, this system of preferences is an ordered logit. Because the likelihood estimation (see Equation (7)) uses data on timing of the first well, the probability that at least one well is drilled is equal to the probability that the firm prefers to drill either one or M wells rather than zero. The ordered logit specification implies that we only need to compare the payoff of drilling zero wells to drilling one well (because if the firm prefers to drill zero wells over one well, it also prefers to drill zero wells over M wells). Therefore, we write the hazard H_{it} as:

$$H_{it} = \frac{\exp(V_{i,t}^1/\sigma_\nu f_{it})}{\exp(V_{i,t}^1/\sigma_\nu f_{it}) + \exp(V_{i,t}^0/\sigma_\nu f_{it})}, \quad \text{(A.18)}$$

and the probability of first drilling at date t as:

$$Pr(I_{it} = 1) = H_{it} \cdot \Pi_{t'=1}^{t-1}(1 - H_{i,t'}) \quad \text{(A.19)}$$

C.5 Restrictions imposed in the sample of units used for maximum likelihood estimation

As we discuss in Section 5.4, we impose restrictions on the sample of units used in the maximum likelihood estimation of our model. We start with the sample of 1226 units in which drilling had not yet occurred by Q1 2009. We first filter out units that had already reached their maximum acreage by this date (38% of units) and units that reach this acreage after 2013 (2% of units), so that all units "start" in-sample. We then drop units in which reported leased acreage ever increases after having reached its maximum (affecting 25% of the original 1226 units) and units that are drilled before reaching their maximum acreage (affecting 30% of the original 1226 units).

We drop a small share (6%) of units with no royalty data. We then mitigate problems with measurement error in reported acreage by dropping units never having more than 160 acres leased in our data (affecting 21% of the original 1226 units) and then re-scaling each unit's leased acreage in each quarter so that the maximum acreage leased in each unit during the sample period is 640 acres (the standard unit size). That is, we multiply the leased acreage in each unit i in each quarter t by 640 and divide by the maximum reported acreage for i during the sample. This re-scaling is motivated by our belief that reported *changes* in relative acreage leased within a unit over time, particularly when changes are driven by lease expiration, are less error-prone that reported *levels* of acreage.

Finally, we drop units in which a well is drilled when leased acreage is very small in our data: less than 10% of the maximum acreage ever leased. This last restriction removes just 3 of the remaining units, following the imposition of the other restrictions. It is necessary

Figure A.6: Locations of units in the estimation sample

because it is extremely difficult for our model to rationalize drilling when acreage is so small (since the firm makes zero profits absent a highly extreme draw of the cost shock ν), but we sometimes see such instances in the data due to missing and mis-reported data in the lease records. The final estimation sample then contains 241 units. We map the distribution of these units within the Haynesville Shale in Figure A.6.

C.6 Maximum likelihood estimation

C.6.1 Integrated likelihood computation

Computing the log likelihood per Equation (7) requires integrating the components of each unit's likelihood function over the distribution $\psi(\theta|\sigma_\theta)$. A natural way to do this would be to use a numerical procedure like quadrature or Monte Carlo integration for each unit, where the evaluation points are functions of the unit's OGIP and the parameters β_0, β_1, and σ_θ. This procedure is computationally intensive, however, because it requires the model to be solved—for each unit—for each evaluation point and for each guess of the parameters in Ω.

We instead adopt a nested loop procedure that minimizes the number of times the model must be solved. The outer-most loop searches over the cost parameters α_0 and σ_ν. At each candidate pair of parameters, the model is solved for each unit on a fixed grid of productivities that captures the range of values of $\beta_0+\beta_1 X_i+\theta_i$ that might plausibly be encountered. This grid contains 200 points, linearly spaced from -45.50 TBtu to -20.50 TBtu. Then an inner loop searches for the parameters β_0, β_1, σ_θ, and σ_ε that maximize the likelihood (conditional on α_0 and σ_ν) by integrating the likelihood for each unit on this grid. This integration occurs via Gauss-Hermite quadrature over 11 nodes, where the node locations (which we interpolate on the grid) are functions of β_0, β_1, and σ_θ, as well as each unit's OGIP value.

C.6.2 Standard error estimation

The standard errors of the estimates obtained from the maximum likelihood procedure account for both spatial correlation of outcomes between nearby sections and the sampling variance of the β_w, γ_0, γ_1, γ_2, and α_1 parameters that are estimated in advance.

First, we address spatial correlation by clustering standard errors at the township level (recall that townships are squares consisting of 36 sections). Let S denote the N by 6 matrix of unit-level likelihood scores, where N is the number of units.[3] Let $A = S'S/N$. Let W denote an N by N matrix with ones in cells corresponding to units in the same township and zeros in all other cells. Let $B = S'WS$. The spatially clustered covariance matrix is then given by $V = A^{-1}BA^{-1}/N^2$.

We then account for the sampling variance of the first stage parameters using the two-step procedure from Murphy and Topel (1985). Let V_1 denote the covariance matrix for β_w, γ_0, γ_1, γ_2, and α_1.[4] Let S_1 denote the N by 5 matrix of likelihood scores with respect to these parameters and let $R = S_1'S/N$. The final covariance estimate for the parameters estimated in the maximum likelihood routine is then given by $V + A^{-1}R'V_1RA^{-1}/N$.

D Additional detail for the model used in Section 7

In Section 7, we present a model in which the owner offers the firm a lease contract in which the bonus is set optimally (in terms of maximizing the owner's expected discounted revenue) given the lease's royalty and primary term. Should the firm accept the contract and then not drill a well during the primary term, the owner offers the firm a renewal—with the same royalty and primary term—in exchange for another optimally-set bonus. This process repeats until the firm either drills a well or decides to not pay the bonus.

This appendix provides more information on how we simulate this model. The computation involves a nested loop. The inner loop solves for each possible firm type's drilling probabilities and value function during a lease term, as a function of its terminal payoff should it not drill a well. The outer loop solves for the owner-optimal bonus (which then determines the firm's payoff at the end of the preceding term) and then iterates an infinite sequence of lease terms until convergence.

The inner loop is the same finite-horizon stopping problem discussed in Section 5.4 and Appendix C above, but with the continuation value upon lease expiration being non-zero for firms that choose to pay the renewal bonus.

We model the renewal bonus as being paid in the final period T of the preceding lease term. For a given bonus value, firms that have not yet drilled compare the bonus against their

[3] We compute all likelihood scores by taking numerical two-sided derivatives.

[4] For β_w, γ_0, γ_1, γ_2, the covariances are computed using the clustered bootstrap procedure discussed in Section 5.1 and Appendix C.3. The estimate of α_1 is assumed to be independent of the other parameters.

expected value from a new lease next period, and they pay the bonus if the latter exceeds the former.[5] The owner then chooses the revenue-maximizing bonus, accounting for both the bonus revenue itself and expected future royalties (and future renewal bonuses) from the types θ that elect to participate. This decision trades off, on the margin, the immediate revenue gain from a higher bonus against the loss of revenue from reduced participation from marginal types.

The outer loop then proceeds as follows, starting from an initial guess of each firm type's continuation value upon expiration:

1. Use finite-horizon backward induction to compute drilling probabilities and the firm's and owner's values during the lease term.

2. Compute the owner-optimal bonus, payable at the end of the preceding term, and compute a new continuation value upon expiration for each firm type, given the bonus.

3. Return to step 1 and iterate until convergence of the firm's and owner's value functions at the start of each new lease term.[6]

We model the initial lease contract as being signed the period before the primary term begins. The owner-optimal bonus value for this contract is then the same as the converged bonus from the above loop.

E Analytic model for the mineral owner's value-maximizing contract

This appendix characterizes analytically a revenue-maximizing take-it-or-leave-it menu of contracts that a mineral owner would offer to a firm with private information on the unit's expected production, where the drilling date and realized production are contractible but completion effort (e.g., water input) is not. There are two main results. First, if the sensitivity of natural gas production to completion effort is small enough, then the mineral owner's revenue-maximizing lease contract involves both a royalty and a provision—here, a drilling subsidy—to accelerate drilling and counteract the royalty's delay incentive. Second, if instead the production function is dominated by the firm's effort choice, the revenue-maximizing contract instead involves a royalty and a drilling tax.

As in the computational model discussed in Section 7, the mineral owner can make a TIOLI offer to a single firm, and then given the contract the firm decides whether and when to drill a well, and if so, how much effort to exert conditional on drilling. We simplify the problem here—thereby gaining analytic tractability—by assuming that only one well can be drilled on the lease and by assuming that the only variable that evolves over time is the natural gas price (therefore abstracting away from rig dayrates and water prices that evolve

[5] Firms realize the period T ν_{iT} shocks before making their period T decision of whether to drill, pay the bonus, or let the lease expire.

[6] The outer loop is not guaranteed to converge but typically does in practice. There are a handful of cases in which the outer loop cycles between values for a few types; in these situations the difference in values is small (less than \$1000) for these types, and we treat these cycles as having converged.

over time, and from the ν_{it} drilling cost shocks that are included in the paper's computational model). The primitives of the model are as follows:

- Time is discrete and denoted by $t \in \{0, ..., T\}$, where T is possibly infinite. The lease contract 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 owner observes the period in which drilling (and completion) occur. Only one well may be drilled on the lease.

- $e \in \mathbb{R}^+$ denotes non-contractible completion "effort" (i.e., water use in the main text).

- The cost of drilling and completing the well is given by $c_0 + c_1 e$, where c_0 and c_1 are strictly positive scalars that are common knowledge.

- If the firm drills, its natural gas production is given by $y = \beta_0 + \theta + g(e) + \varepsilon$.

  - $\beta_0 \in \mathbb{R}$ is common knowledge. $\theta \in \mathbb{R}$ is known by the firm but not by the owner. $\Psi(\theta)$ denotes the owner's rational belief about the distribution of θ. The expected value of θ is 0, and $\Psi(\theta)$ has support on $[\theta_L, \bar{\theta}]$.
  - The function $g(e)$ is common knowledge and maps completion effort onto gas production, with the properties that $g'(e) > 0$ and $g''(e) < 0$ $\forall e$, and that $\lim_{e\to 0^+} g'(e) \to \infty$ and $\lim_{e\to\infty} g'(e) \to 0$ (we will later specify $g(e) = \beta_w \log e$ as in our computational model).
  - ε is a mean-zero disturbance that is unknown by the owner and firm prior to drilling and the choice of e. ε is orthogonal to e and θ, and its distribution function $\Lambda(\varepsilon)$ is common knowledge.
  - Output y is contractible, and for simplicity we assume that y is completely realized in the same period that the well is drilled and completed.

- The gas price at time t is denoted P_t and is common knowledge. The gas 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.

Both the owner and firm are risk neutral, share a common per-period discount factor δ, and seek to maximize the expected present value of their respective cash flows. At $t = 0$, the owner 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).

We first characterize the firm's problem and then turn to the owner's contract design problem. The characterization closely follows parts of Laffont and Tirole (1986) and Board (2007). 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 and the price at execution will be necessary for optimality.

E.1 Firm's problem

The firm must make three decisions, in sequence:

1. Report a type $\hat{\theta}$ to the owner at $t = 0$ (or opt-out)
2. Choose a time $\tau \in \{1, ...T\}$ 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 (suppressing the dependence of z on $\hat{\theta}$, y, and P^t). The firm's problem, conditional on participation, is then given by

$$\max_{\hat{\theta}, \tau^*(\theta,z), e} \Pi(\hat{\theta}, \tau^*(\theta, z), e, \theta) = E_P[(P_\tau(\beta_0 + \theta + g(e)) - (c_0 + c_1 e) - z_\tau(\hat{\theta}, \beta_0 + \theta + g(e) + \varepsilon, P^t))\delta^\tau] - R(\hat{\theta}), \tag{A.20}$$

where E_P is the expectation at the start of period 0, taken over all prices. Note that total surplus is maximized by the solution to (A.20) when the transfers z and R are set to zero.

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.[7]

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 owner. For truth-telling to be incentive compatible, it must be the case that $e(\hat{\theta})$ 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 our attention to deviations in a *concealment set* in which, for any report $\hat{\theta}$, the chosen effort $\tilde{e}$ is such that $\theta + g(\tilde{e}) = \hat{\theta} + g(e(\hat{\theta}))$. Thus, absent uncertainty generated by ε, any deviation outside the concealment set can be detected by the owner.[8]

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 $g(H(E)) = E$. The derivatives of g and H are related by $H'(g(e)) = 1/g'(e)$. The firm's problem may then be written:

$$\max_{\hat{\theta}} \Pi(\hat{\theta}, \theta) = E_P[(P_{\tau(\hat{\theta})}(\beta_0 + \hat{\theta} + g(e(\hat{\theta}))) - (c_0 + c_1 H(\hat{\theta} - \theta + g(e(\hat{\theta})))) - z_{\tau(\hat{\theta})}(\hat{\theta}, \beta_0 + \hat{\theta} + g(e(\hat{\theta})) + \varepsilon, P^{\tau(\hat{\theta})}))\delta^{\tau(\hat{\theta})}] - R(\hat{\theta}). \tag{A.21}$$

[7] Board (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 assumption that P_t is bounded above replacing the Board (2007) assumption that costs are bounded below.

[8] In the presence of a non-degenerate distribution $\Lambda(\varepsilon)$, the sufficiency conditions for implementing the mechanism will be stricter than the conditions given below that the optimal stopping time is decreasing in θ, and $\theta + g(e(\theta))$ is increasing in θ. In the event that they are not satisfied, the owner will need to "iron" over regions in the type space where incentive compatibility does not hold.

To obtain the marginal information rent for a firm of type θ, we use the generalized envelope theorem from Milgrom and Segal (2002) and take the partial derivative of $\Pi(\hat{\theta}, \theta)$ with respect to θ:

$$\left.\frac{\partial \Pi(\hat{\theta}, \theta)}{\partial \theta}\right|_{\hat{\theta}=\theta} = E_P\left[\frac{c_1 \delta^{\tau(\theta)}}{g'(e(\theta))}\right]. \tag{A.22}$$

Equation (A.22) is the first-order incentive compatibility condition. The second order monotonicity condition is

$$\frac{\partial \Pi(\hat{\theta}, \theta)}{\partial \theta \partial \hat{\theta}} \geq 0. \tag{A.23}$$

From taking derivatives of Equation (A.21), the mechanism will satisfy condition (A.23) if the optimal stopping time is decreasing in $\hat{\theta}$, and $\hat{\theta} + g(e(\hat{\theta}))$ is increasing in $\hat{\theta}$.[9]

Given incentive compatibility, integration of Equation (A.22) yields the firm's information rent:

$$\Pi(\theta, \theta) = E_P\left[\int_{\underline{\theta}}^{\theta} \frac{c_1 \delta^{\tau(s)}}{g'(e(s))} ds\right], \tag{A.24}$$

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

E.2 Revenue-maximizing contract for the owner

Continuing to follow Laffont and Tirole (1986) and Board (2007), we treat the owner'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 total surplus minus information rent is maximized. We therefore write the owner's problem as

$$\max_{\tau(\theta), e(\theta), \underline{\theta}} \int_{\underline{\theta}}^{\bar{\theta}} \left[E_P\left(P_{\tau(\theta)}(\beta_0 + \theta + g(e(\theta))) - c_0 - c_1 e(\theta)\right)\delta^{\tau(\theta)} - \int_{\underline{\theta}}^{\theta} \frac{c_1 \delta^{\tau(s)}}{g'(e(s))} ds\right] \psi(\theta) d\theta, \tag{A.25}$$

where the owner 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 $h(\theta) \equiv f(\theta)/(1 - F(\theta))$ denote the hazard function, we re-write the owner's problem as:

$$\max_{\tau(\theta), e(\theta), \underline{\theta}} \int_{\underline{\theta}}^{\bar{\theta}} E_P\left[\left(P_{\tau(\theta)}(\beta_0 + \theta + g(e(\theta))) - c_0 - c_1 e(\theta) - \frac{c_1}{h(\theta) g'(e(\theta))}\right)\delta^{\tau(\theta)}\right] \psi(\theta) d\theta. \tag{A.26}$$

Now recall the firm's problem, Equation (A.20). Following the logic in Board (2007), the owner can induce the firm to follow the stopping rule implied by (A.26) by setting the contingent payment z equal to the information rent term in (A.26), since doing so makes the

[9] The owner's optimal timing and effort functions defined below in equations (A.26) and (A.28) will satisfy the condition that the optimal stopping time is decreasing in θ if $h'(\theta) \geq 0$. To see this sufficiency, first observe that $h'(\theta) \geq 0$ is sufficient for the total derivative of the term in parentheses in (A.26) to be 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. The second condition—that $\theta + g(e(\theta))$ is increasing in θ—is difficult to characterize in terms of primitives.

firm's problem equivalent to the owner's problem. Thus, the revenue-maximizing contingent payment is given by:

$$z_\tau(\theta, y, P^t) = \frac{c_1}{h(\theta)g'(e(\theta))}. \tag{A.27}$$

The contingent payment in (A.27) is positive, which will lead to delayed drilling relative to the social optimum. Note that the payment is zero for the highest type $\bar{\theta}$ firm because $1/h(\bar{\theta}) = 0$, reflecting the standard "no distortion at the top" rule. The optimal up-front payment $R(\theta)$ is set to equate the firm's payoff to the information rent expressed in Equation (A.24), where the payoff of the endogenously chosen type $\underline{\theta}$ firm is set to zero.

The 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. Because effort is not contractible, this mechanism must be implemented using a payment that is contingent on production y. Paralleling Laffont and Tirole (1986), we 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 above, since the affine payment structure constrains punishments for deviations outside of the concealment set.

To derive the optimal linear production tax, we first take the pointwise derivative of (A.26) with respect to $e(\theta)$ to obtain the FOC that defines the owner's optimal effort function, conditional on drilling at τ. Suppressing the dependence of $g(e(\theta))$ and its derivatives on θ, this FOC is given by:

$$FOC_{e(\theta)} : P_\tau g' - c_1 + \frac{c_1 g''}{h(\theta)g'^2} = 0. \tag{A.28}$$

Because $g'' < 0$, FOC (A.28) implies that $e(\theta)$ must be strictly less than the surplus-maximizing effort, except for type $\bar{\theta}$.

To obtain the optimal production tax, we return to the firm's problem and take the derivative of Equation (A.20) with respect to e to derive the firm's FOC for its optimal effort, conditional on drilling at τ:

$$FOC_e : P_\tau g' - c_1 - \frac{\partial z_\tau(\theta, y, P^t)}{\partial y} g' = 0 \tag{A.29}$$

Combining equations (A.28) and (A.29) yields the linear tax on production that aligns the firm's incentives with the effort function that the owner wishes to induce:

$$\frac{\partial z_\tau(\theta, y, P^t)}{\partial y} = \frac{-c_1 g''}{h(\theta)g'^3} \tag{A.30}$$

We can now solve for the optimal contingent payment using equations (A.27) and (A.30):

$$z_\tau(\theta, y, P^t) = \frac{c_1}{h(\theta)g'} - \frac{c_1 g''}{h(\theta)g'^3}(y - \beta_0 - \theta - g) \tag{A.31}$$

Rearranging and using Equation (A.29) to eliminate c_1, we obtain:

$$z_\tau(\theta, P^t, y) = \frac{P_\tau(g'^2 + g''(\beta_0 + \theta + g))}{h(\theta)g'^2 - g''} - \frac{g''}{h(\theta)g'^2 - g''} P_\tau y \tag{A.32}$$

The second term in (A.32) is a positive tax on revenue $P_\tau y$; i.e., a royalty. The first term is a transfer at the time of drilling that is not dependent on output. It may be positive (i.e., a tax on drilling) or negative (a drilling subsidy).

E.3 Drilling subsidy or drilling tax?

The sign of the first term of Equation (A.32) depends on the function $g(e)$. We now tie our analysis in this section more closely to the main text and adopt the functional form $g(e) = \beta_w \log e$, where $\beta_w \in \mathbb{R}^{++}$. With this functional form assumption, we can rewrite Equation (A.32) as:

$$z_\tau(\theta, P^t, y) = \frac{P_\tau(\beta_w - \beta_0 - \theta - \beta_w \log e)}{1 + h(\theta)\beta_w} + \frac{1}{1 + h(\theta)\beta_w} P_\tau y \tag{A.33}$$

The numerator of the first term in Equation (A.33) may be positive or negative, and the denominator is guaranteed to be strictly positive. To better understand the sign of this term, we now consider a comparative static in which we change the sensitivity of output y to effort e, holding expected production and effort fixed for the mean type $\theta = 0$. More precisely, we introduce a scalar $b < \beta_w$ that adjusts the marginal productivity of effort so that:

- The production function is $y = (\beta_0 + b \log e_0^*) + \theta + (\beta_w - b) \log e + \varepsilon$, where e_0^* denotes the surplus-maximizing effort level of the mean $\theta = 0$ type under a non-distortionary contract, at the initial condition of $b = 0$.
- The cost of drilling is $c_0 + c_1 \frac{(\beta_w - b)}{\beta_w} e$.

The addition of $b \log e_0^*$ to β_0 and the multiplication of c_1 by $(\beta_w - b)/\beta_w$ ensure that after β_w is adjusted by subtracting b, then the mean firm type, in the absence of a distortionary contract, would choose the same effort e_0^* and obtain the same expected production, conditional on drilling at the same trigger price, as was the case under $b = 0$. But changes in effort will now have a reduced impact on expected output if $b \in (0, \beta_w)$ and a greater impact if $b < 0$.

The numerator of Equation (A.33) is now given by the expression:

$$P_\tau(\beta_w - b - \beta_0 - b \log e_0^* - \theta - \beta_w \log e_\theta + b \log e_\theta), \tag{A.34}$$

where we now write e_θ rather than just e to clarify that this value represents the effort level of each participating type θ under the mechanism. e_θ^* denotes the surplus-maximizing effort of each type θ.

Now consider taking $b \to \beta_w$, making production less sensitive to effort. There will be a value of b sufficiently close to β_w such that the sign of (A.34) will be determined by the sign of $-\beta_0 - \theta - b \log e_0^*$. This expression must be strictly negative for type $\theta = 0$ if it participates, since as $b \to \beta_w$ it approaches the negative of that type's expected production conditional on drilling under a non-distortionary contract, which cannot be negative if it participates. In that case, the expression will also be strictly negative for all types $\theta > 0$ as well. If the $\theta = 0$ type does not participate, then the expression will hold for all participating $\theta > 0$ types, since for those types $e_0^* > e_\theta^*$ (because in a non-distorted contract, higher types exert less effort) and because $-\beta_0 - \theta - b \log e_\theta^* < 0$ (since those types participate).

To complete the proof, we need to consider the possibility that types $\theta < 0$ may participate. In this case, it is sufficient to ensure that b is close enough to β_w that $e_\theta \leq e_0^*$ for all such types. Such a value of b is guaranteed to exist, since the royalty increases with b and drives effort towards zero as $b \to \beta_w$ for all but the highest type under the optimal contract. From there, since $-\beta_0 - \underline{\theta} - b \log e_{\underline{\theta}} < 0$ for the lowest participating type, it follows that $-\beta_0 - \underline{\theta} - b \log e_0^* < 0$ for that type and all higher types.

Finally, consider the opposite comparative static in which the production function becomes increasingly sensitive to effort by evaluating the case of $b < 0$. For a sufficiently negative b, expression (A.34) becomes dominated by the term $-b - b \log e_0^* - \theta + b \log e_\theta$. This term is guaranteed to be positive $\forall \theta$ under the sufficient condition from subsection E.1 that $\theta + g(e(\theta))$ is increasing in θ. First, observe that for $\theta \geq 0$, we have $e_0^* \geq e_\theta^* > e_\theta$. From there, making b sufficiently negative that $-b > \bar{\theta}$ is sufficient for $-b - b \log e_0^* - \theta + b \log e_\theta$ to be strictly positive $\forall \theta \geq 0$. Second, the condition that $\theta + g(e(\theta))$ is increasing in θ implies that, for negative enough b, $-\theta + b \log e_\theta$ is decreasing in θ. The fact that $-b - b \log e_0^* - \theta + b \log e_\theta$ is strictly positive for $\theta = 0$ then implies that the expression is strictly positive $\forall \theta < 0$.

Thus, if the production function is dominated by the firm's choice of effort, the owner's revenue-maximizing contract involves a drilling tax rather than a drilling subsidy. The model and result in this case are actually similar to that of Board (2007): the high sensitivity of output to input choice makes it difficult to contract on output (Board (2007) completely rules out contracting on output), so that the owner's revenue-optimal contract then involves a tax on exercising the drilling option instead (as in Board (2007)).

We have quantitatively examined the optimal fixed payment using our computational model of Section 7. We find that at our estimate of β_w and with a 25% royalty, the owner would maximize expected revenue with a drilling tax of \$0.34 million. However, when $\beta_w = 0$ and the royalty is 39%, the owner would maximize expected revenue with a drilling subsidy of \$1.26 million, consistent with the analytic model above.

References

Ackerberg, Daniel A. and Paul J. Devereux, "Improved JIVE Estimators for Overidentified Linear Models with and without Heteroskedasticity," *The Review of Economics and Statistics*, 2009, pp. 351–362.

Anderson, Soren T., Ryan Kellogg, and Stephen W. Salant, "Hotelling Under Pressure," *Journal of Political Economy*, June 2018, *126*, 984–1026.

Bloomberg, "U.S. Crude oil and natural gas futures prices," accessed 3 October, 2017.

Board, Simon, "Selling Options," *Journal of Economic Theory*, 2007, *136*, 324–340.

Bureau of Labor Statistics, "Consumer Price Index, all urban, all items less energy, not seasonally adjusted (CUUR0000SA0LE)," https://www.bls.gov/cpi/data.htm; accessed 8 July, 2018.

Chetty, Raj, John N. Friedman, Tore Olsen, and Luigi Pistaferri, "Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records," *Quarterly Journal of Economics*, 2011, *126*, 749–804.

Enverus, "DI Classic completion, lease, and well data," accessed in batches on 1 June, 2012, 6 May, 2014, 7 January, 2016, 21 January 2016, and 22 February, 2016.

_ , "DI Desktop monthly well production data," accessed on 5 February, 2016.

_ , "Rig Day Rate Report," accessed on 2 November, 2017.

Gülen, Gürcan, Svetlana Ikonnikova, John Browning, Katie Smye, and Scott W. Tinker, "Production Scenarios for the Haynesville Shale Play," *SPE Economics & Management*, 2015, *7* (5).

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

Kolesàr, Michal, "Estimation in an Instrumental Variables Model with Treatment Effect Heterogeneity," 2013. Working paper.

Laffont, Jean-Jacques and Jean Tirole, "Using Cost Observation to Regulate Firms," *Journal of Political Economy*, Jun. 1986, *94* (3), 614–641.

Louisiana Department of Natural Resources, "Well and unit shapefiles," http://sonris-www.dnr.state.la.us/gis/agsweb/IE/JSViewer/index.html?TemplateID=181; accessed 7 January, 2016.

_ , "Well drilling and completion cost reports," http://sonris.com/; accessed in spring and summer, 2016.

_ , "Well tabular data," http://sonlite.dnr.state.la.us/pls/apex/f?p=108:2:32690097475501:::::; accessed in May, 2016.

Male, Frank, Akand W. Islam, Tad W. Patzek, Svetlana Ikonnikova, John Browning, and Michael P. Marder, "Analysis of Gas Production from Hydraulically Fractured Wells in the Haynesville Shale Using Scaling Methods," *Journal of Unconventional Oil and Gas Resources*, 2015, *10*, 11–17.

Milgrom, Paul R. and Ilya Segal, "Envelope Theorems for Arbitrary Choice Sets," *Econometrica*, Mar. 2002, *70* (2), 583–601.

Murphy, Kevin M. and Robert H. Topel, "Estimation and Inference in Two-Step Econometric Models," *Journal of Business & Economic Statistics*, 1985, *3* (4), 370–379.

Patzek, Tad W., Frank Male, and Michael Marder, "Gas Production in the Barnett Shale Obeys a Simple Scaling Theory," *PNAS*, Dec. 2013, *110* (49), 19731–19736.

Robinson, Peter M., "Root-n-consistent Semiparametric Regression," *Econometrica*, Jul. 1988, *56* (4), 931–954.

US Census Bureau, "County Boundary Files for Louisiana," https://www2.census.gov/geo/pvs/tiger2010st/22_Louisiana/22/tl_2010_22_county10.zip; accessed 28 April, 2020.