Econometrica Supplementary Material

SUPPLEMENT TO "DRILLING DEADLINES AND OIL AND GAS DEVELOPMENT"

(*Econometrica*, Vol. 92, No. 1, January 2024, 29–60)

EVAN HERRNSTADT
Congressional Budget Office

RYAN KELLOGG
Harris School of Public Policy, University of Chicago and National Bureau of Economic Research

ERIC LEWIS
Bush School of Government and Public Service, Texas A&M University

APPENDIX 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 (2012).
- Enverus production data from Enverus (2016). 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 (2012). Enverus collects data on leases signed in Louisiana. Further details are below.
- Enverus dayrate data from Anderson, Kellogg, and Salant (2018) and Enverus (2017). We use dayrates that correspond to the "ArkLaTx" region, for rigs with depth

Evan Herrnstadt: evan.herrnstadt@cbo.gov
Ryan Kellogg: kelloggr@uchicago.edu
Eric Lewis: ericlewis@tamu.edu

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, Ikonnikova, Browning, Smye, and Tinker (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, Male, and Marder (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:

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

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, Male, and Marder (2013), we limit the sample to observations that are the fourth month or later ($t \ge 4$) 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, Islam, Patzek, Ikonnikova, Browning, and Marder (2015) for the Haynesville, which finds that τ is 14.16 months. We then use nonlinear 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):

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

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:

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

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 released 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, that is, 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:

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

Here, k indexes characteristics of the observation, for example, 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.[A1] 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 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 second reason is that the well actually targets multiple units. In cases where a wellbore passes through multiple units, we only match a well to a unit if at least 300 meters of the horizontal wellbore pass through the unit.

[A1] 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.

APPENDIX 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, Friedman, Olsen, and Pistaferri (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).

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).

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: This 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.

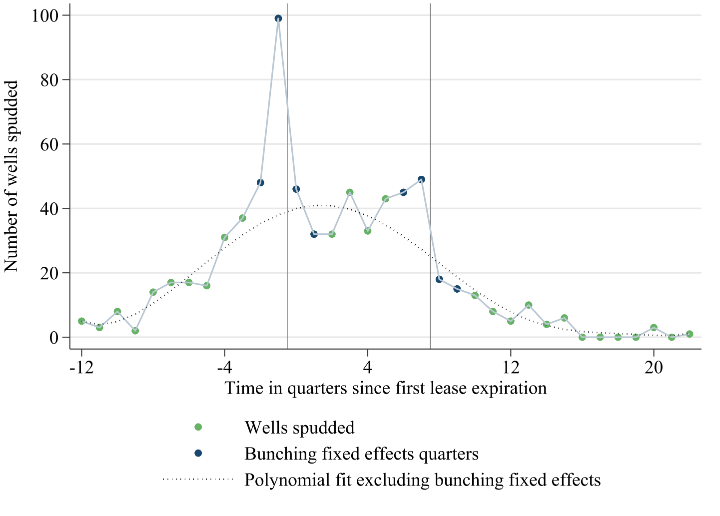

FIGURE A.1.—Estimates from bunching analysis. Note: The figure presents data and estimates corresponding to the bunching analysis in column (1) of Table A.I. Plotted data include counts of wells spudded, the quarters to which we add bunching fixed effects, and the polynomial predicted probabilities given 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 2 years after first lease expiration.

B.2. *Additional Descriptive Figures and Table*

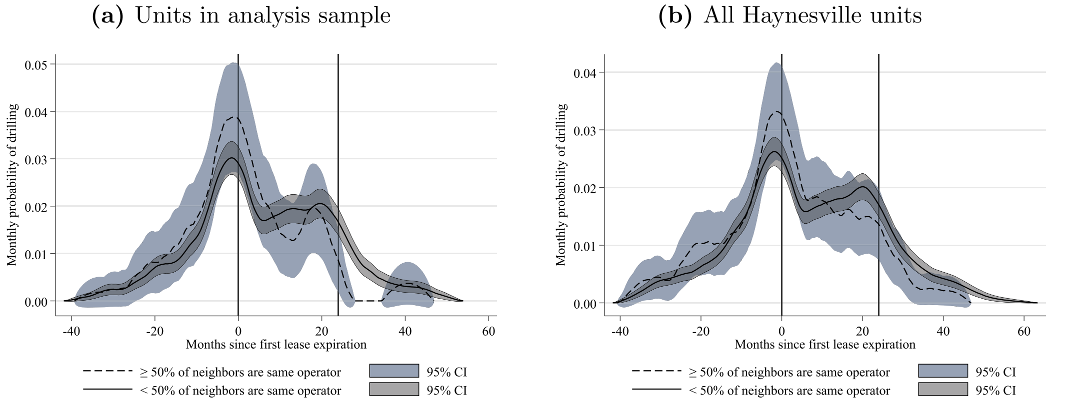

FIGURE A.2.—Comparison of units with identical vs. different neighboring operators. Note: The figure shows kernel-smoothed estimates of the probability of drilling the first Haynesville well in a unit on a given date, relative to the expiration date of the first lease within the unit to expire. Vertical lines are drawn at the date of first lease expiration and 2 years after first lease expiration. Figure compares units in which ≥ 50% of the nearby units have the same operator vs. units where ≥ 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 Section 3.4. In Panel (b), we show results using all Haynesville units.

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	29,370	37,528	37,448

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, for example, 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.

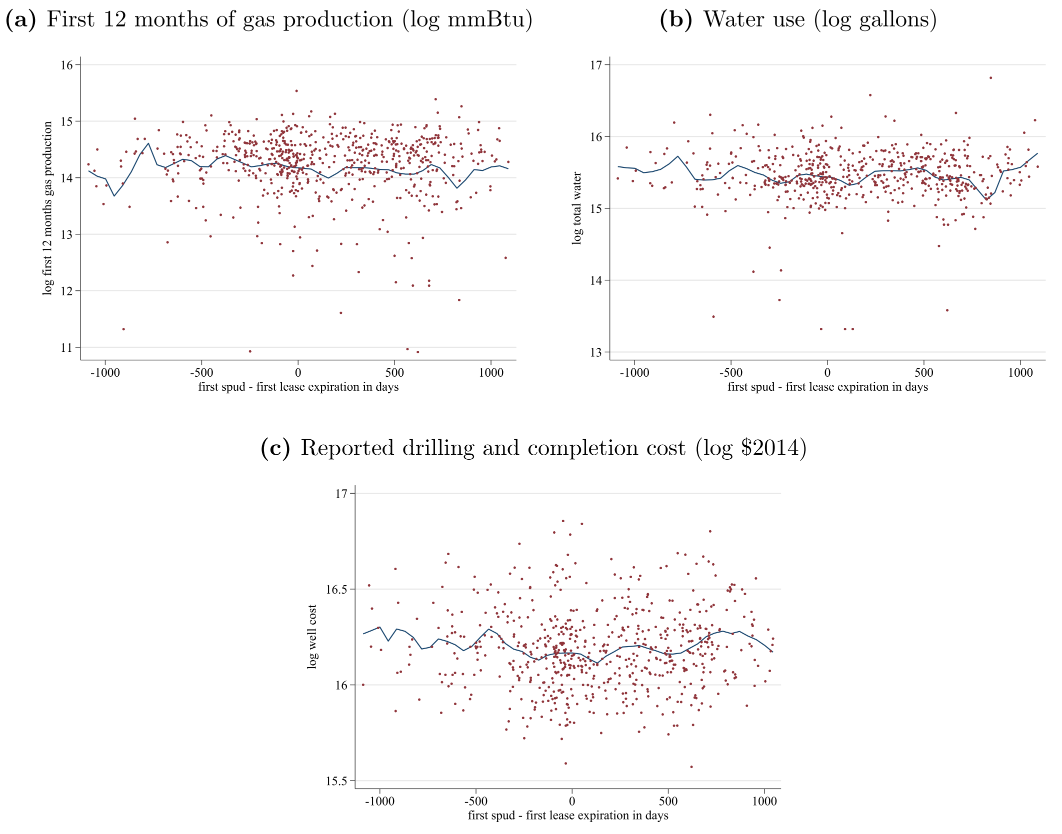

FIGURE A.3.—Wells' production, water use, and cost vs. time relative to first lease expiration. Note: The figure plots natural gas production, water input, and reported drilling cost for the first well drilled in each unit against the well's spud date relative to the date of first lease expiration (measured in days). The line is the predicted value from a local polynomial regression.

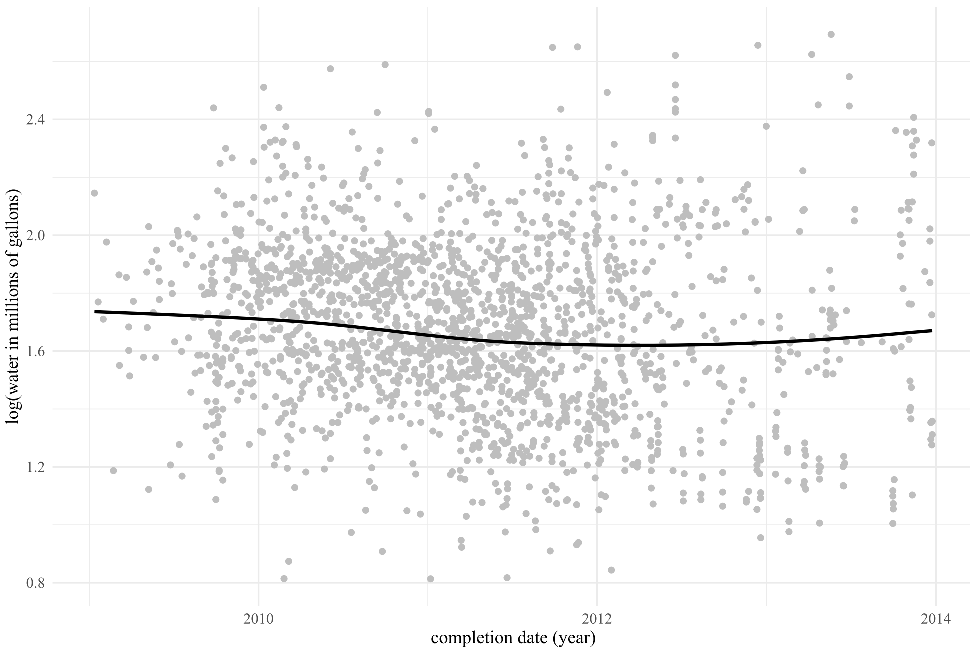

FIGURE A.4.—Log water use versus time. Note: The figure presents a scatter plot and lowess estimate of each well's log water use versus its spud date. Scatter plot excludes a small number of outlier water observations.

APPENDIX 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 nonparametric 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:

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

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 $\sigma_\phi = 3632$ meters. If we instead use production rather than log water as an input to calculate σ_ϕ, we calculate $\sigma_\phi = 2407$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 nonparametric 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 nonparametric 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 nonparametric 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 nonparametric 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 nonparametric 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 nonparametric 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 nonparametric 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.

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:[A2]

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

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 5000 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

[A2]We use the first-order condition for when the well "pays out" (as opposed to equation (A.9)).

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	2019

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

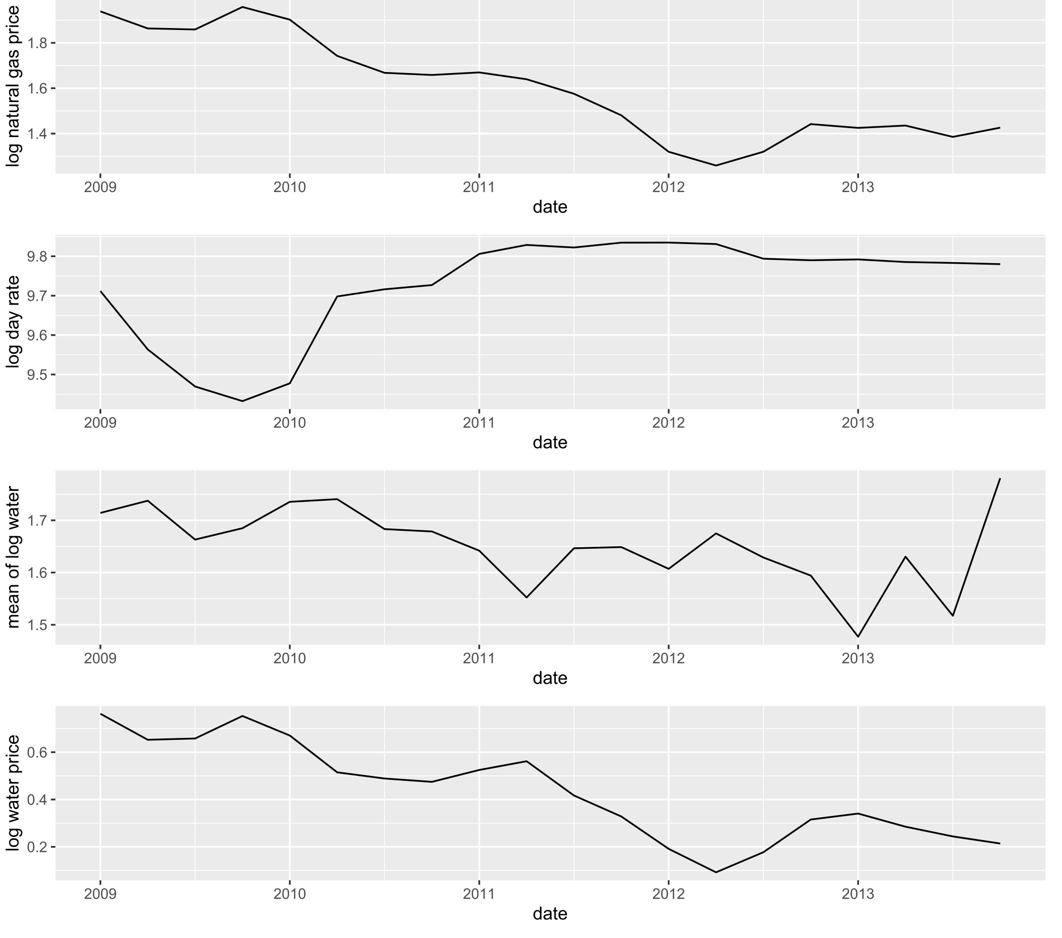

FIGURE A.5.—Estimated log water price relative to natural gas price, dayrate, and average log water use. Note: The natural gas price is in \$/mmBtu, the dayrate is in \$/day, water use is in millions of gallons, and the water price is in \$/gallon.

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.

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}:

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

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:

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

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:

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

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

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

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

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

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

$$\begin{aligned}\nu_{it}^1 - \nu_{it}^0 > \frac{1}{f_{it}}\big[&M\delta E[\tilde{V}_{i,t+1}|S_{it}, \theta_i] - (M-1)\delta E[\tilde{U}_{i,t+1,t,\theta}|S_{it}, \theta_i] \\ &- E_{P_{wt}}\big[\pi_{it}(S_{it}, \theta_i)\big]\big].\end{aligned} \tag{A.17}$$

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

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

and the probability of first drilling at date t as

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

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 rescaling

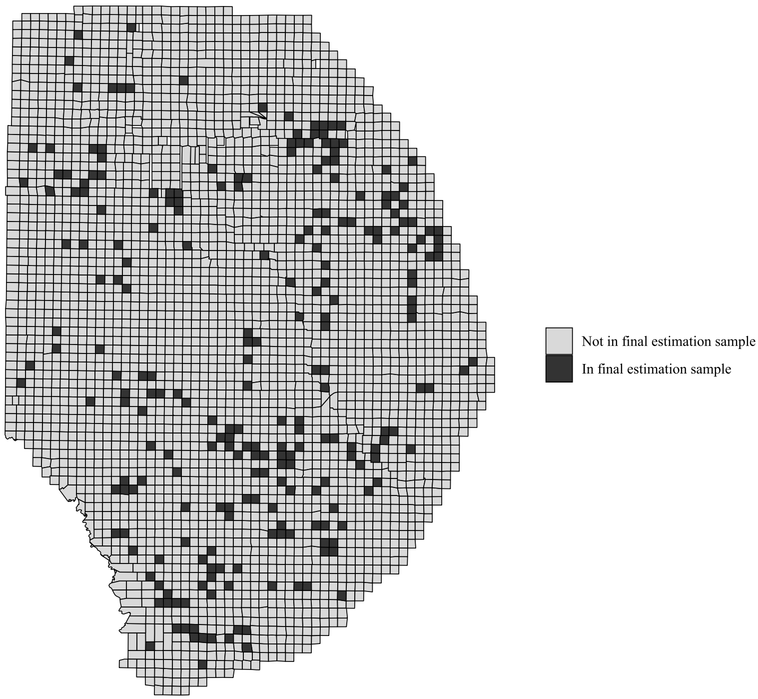

FIGURE A.6.—Locations of units in the estimation sample.

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 rescaling 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 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 misreported 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 -45.50TBtu to -20.50TBtu. 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 11nodes, 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.[A3] 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.[A4] 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 $V + A^{-1}R'V_1RA^{-1}/N$.

APPENDIX 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 nonzero for firms that choose to pay the renewal bonus.

[A3] We compute all likelihood scores by taking numerical two-sided derivatives.

[A4] 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.

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 expected value from a new lease next period, and they pay the bonus if the latter exceeds the former.[A5] 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.[A6]

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.

APPENDIX 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 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 $t \in \{0, \dots, T\}$, 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

[A5] 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.

[A6] 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.

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 noncontractible 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 $\tau \in \{1, \dots T\}$ at which to exercise the option to drill, where $\tau = \infty$ signifies not drilling.
3. Conditional on drilling, select an effort level $e \in \mathbb{R}^+$.

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

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

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.[A7]

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.[A8]

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

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

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:

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

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 $\hat{\theta}$.[A9]

[A7] 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.

[A8] In the presence of a nondegenerate 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.

[A9] 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.

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

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

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

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

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 rewrite the owner's problem as

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

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

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

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

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

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:

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

Combining equations (A.28) and (A.29) yield 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 $P_\tau y$; that is, 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 nondistortionary 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 nondistortionary 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 nondistorted 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 $e_\theta \le e_0^*$ for all such types. Such a value of b is guaranteed to exist, since the royalty increases with b and drives effort toward 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 Section E.1 that $\theta + g(e(\theta))$ is increasing in θ. First, observe that for $\theta \ge 0$, we have $e_0^* \ge e_\theta^* > e_\theta$. 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 (2009): "Improved JIVE Estimators for Overidentified Linear Models With and Without Heteroskedasticity," The Review of Economics and Statistics, 351–362. [11]

ANDERSON, SOREN T., RYAN KELLOGG, AND STEPHEN W. SALANT (2018): "Hotelling Under Pressure," *Journal of Political Economy*, 126, 984–1026. [1]

BLOOMBERG (2017): "U.S. Crude Oil and Natural Gas Futures Prices," accessed 3 October, 2017. [2]

BOARD, SIMON (2007): "Selling Options," *Journal of Economic Theory*, 136, 324–340. [19-21,24]

BUREAU OF LABOR STATISTICS (2018): "Consumer Price Index, All Urban, All Items Less Energy, not Seasonally Adjusted (CUUR0000SA0LE)," https://www.bls.gov/cpi/data.htm; accessed 8 July, 2018. [2]

CHETTY, RAJ, JOHN N. FRIEDMAN, TORE OLSEN, AND LUIGI PISTAFERRI (2011): "Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence From Danish Tax Records," *Quarterly Journal of Economics*, 126, 749–804. [6]

ENVERUS (2012): "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. [1]

——— (2016): "DI Desktop Monthly Well Production Data," accessed on 5 February, 2016. [1]

——— (2017): "Rig Day Rate Report," accessed on 2. [1]

GÜLEN, GÜRCAN, SVETLANA IKONNIKOVA, JOHN BROWNING, KATIE SMYE, AND SCOTT W. TINKER (2015): "Production Scenarios for the Haynesville Shale Play," SPE Economics & Management, 7 (5). [2,3]

KELLOGG, RYAN (2014): "The Effect of Uncertainty on Investment: Evidence From Texas Oil Drilling," *American Economic Review*, 104 (6), 1698–1734. [3]

KOLESÀR, MICHAL (2013): "Estimation in an Instrumental Variables Model With Treatment Effect Heterogeneity," Working paper. [11]

LAFFONT, JEAN-JACQUES, AND JEAN TIROLE (1986): "Using Cost Observation to Regulate Firms," *Journal of Political Economy*, 94 (3), 614–641. [19-21]

LOUISIANA DEPARTMENT OF NATURAL RESOURCES (2016a): "Well and Unit Shapefiles," accessed 7, http://sonris-www.dnr.state.la.us/gis. [1]

LOUISIANA DEPARTMENT OF NATURAL RESOURCES (2016b): "Well Drilling and Completion Cost Reports," accessed in spring and summer, http://sonris.com/. [1]

LOUISIANA DEPARTMENT OF NATURAL RESOURCES (2016c): "Well Tabular Data," accessed in May, 2016 http://sonlite.dnr.state.la.us/pls/apex/f?p=108:2:32690097475501. [1]

MALE, FRANK, AKAND W. ISLAM, TAD W. PATZEK, SVETLANA IKONNIKOVA, JOHN BROWNING, AND MICHAEL P. MARDER (2015): "Analysis of Gas Production From Hydraulically Fractured Wells in the Haynesville Shale Using Scaling Methods," *Journal of Unconventional Oil and Gas Resources*, 10, 11–17. [3]

MILGROM, PAUL R., AND ILYA SEGAL (2002): "Envelope Theorems for Arbitrary Choice Sets," *Econometrica*, 70 (2), 583–601. [20]

MURPHY, KEVIN M., AND ROBERT H. TOPEL (1985): "Estimation and Inference in Two-Step Econometric Models," *Journal of Business & Economic Statistics*, 3 (4), 370–379. [17]

PATZEK, TAD W., FRANK MALE, AND MICHAEL MARDER (2013): "Gas Production in the Barnett Shale Obeys a Simple Scaling Theory," *PNAS*, 110 (49), 19731–19736. [2]
ROBINSON, PETER M. (1988): "Root-n-Consistent Semiparametric Regression," *Econometrica*, 56 (4), 931–954. [11]
US CENSUS BUREAU (2020): "County Boundary Files for Louisiana," https://www.census.gov/geographies/mapping-files/time-series/geo/cartographic-boundary.html, accessed 28 April, 2020. [1]

Co-editors Aviv Nevo and Guido W. Imbens handled this manuscript.

Manuscript received 6 May, 2020; final version accepted 19 November, 2023; available online 27 November, 2023.