Investment, Productivity, and Selection in the U.S. Shale Boom

Thomas R. Covert Konan Hara
Ryan Kellogg Richard L. Sweeney*

March 23, 2026

Preliminary draft — do not cite without permission

Abstract

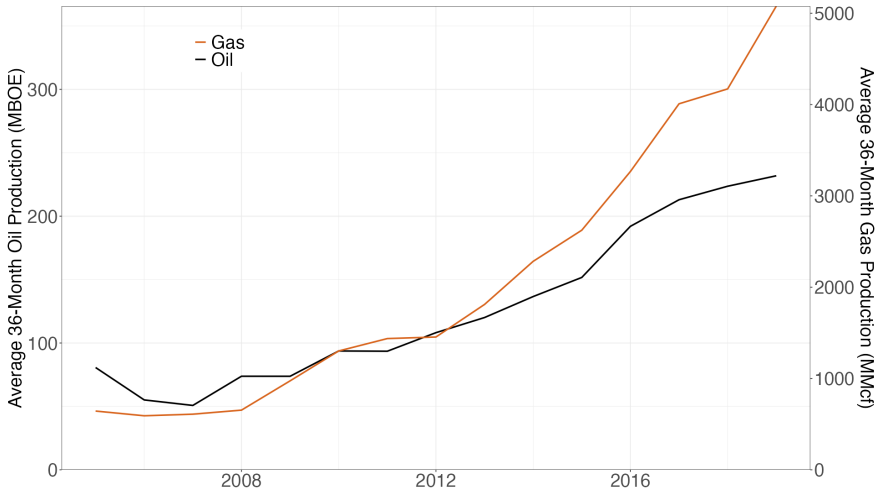
The U.S. shale oil and gas boom has driven an unprecedented increase in U.S. oil and gas production. The primary driver of this ongoing boom has been a large increase in production per new well, rather than increases in the rate of drilling new wells. The goal of this paper is to understand why output per well has increased so dramatically. A core question is the extent to which production growth arose from changes in well site selection—i.e., changes in the underlying geological quality of the precise areas being drilled—or from changes in firms' operational decisions, such as changes in inputs and adoption of new technologies. We develop a joint model of well-level production and drilling decisions that allows us to address site selection that comes from firms' private ex ante signals of sites' underlying geologic quality and from firms' ability to learn about quality over time by observing drilled wells' production outcomes. The model is designed to account for not just whether but when sites are drilled. We estimate the model using data on drilling, production, and leasing from two major U.S. shale plays: the Bakken in North Dakota and the Haynesville in Louisiana. We find that site selection has been important: firms initially drilled areas for which they had initial signals of high quality, and then they later repeatedly drilled in areas that had positive production outcomes. The effects of this site selection were swamped, however, by large increases in output that are associated with increases in input use, driven by either adoption of complementary technology or firms' learning about how to use inputs and existing technology more effectively. The production improvements occur within firms, rather than being driven by reallocation of drilling activity to more productive firms.

*Covert: AirBnb; thom.covert@gmail.com. Hara: Department of Economics, Michigan State University; harakon1@msu.edu. Kellogg: Harris School of Public Policy, University of Chicago; and National Bureau of Economic Research; kelloggr@uchicago.edu. Sweeney: Boston College Department of Economics; sweeneri@bc.edu. This project was supported by funding from NSF award #1919499. We thank Michael Cahana, Pranjal Drall, Madeleine Higgins, Simoni Jain, Alex Lan, and Lixi Wang for valuable research assistance.

1 Introduction

The shale oil and gas revolution has been responsible for an unprecedented increase in U.S. hydrocarbon production. The magnitude of this surge in production was summed up recently by the Energy Information Administration, which noted that the United States now "produces more crude oil than any country, ever" (U.S. Energy Information Administration, 2024). This development has reshaped global markets for oil and natural gas, impacted the trajectory of CO2 emissions, and arguably delayed the transition to zero-emission energy sources (Acemoglu, Aghion, Barrage and Hémous, 2023). An important feature of this surge in aggregate production is that it is not simply coming from drilling more and more wells. Rather, output per new well rose dramatically over time, with wells drilled in 2019 producing more than four times as much as wells drilled a decade earlier (figure 1).

Figure 1: Ouput per new shale well by vintage



Notes: Output from the first 36 months of production for wells drilled in each year, across all U.S. shale plays. See section 2 for details on data sources and construction.

The goal of this paper is to decompose changes in output per well into changes in site selection—i.e., changes in the underlying geological quality of the precise areas being drilled—and changes in firms' operational decisions, such as changes in inputs and the adoption of new technologies. Beyond the intrinsic value of understanding the drivers of one of the most profound developments in energy markets in this century, this decomposition also sheds light on the future of shale. Because the supply of high-quality well locations is finite, the capacity of the site-selection mechanism to sustain long-run production improvements is inherently limited. In contrast, if production gains primarily arose from technology or input choices conditional on location selection—and especially if such improvements came about despite decreases over time in the quality of wells' locations—then one could imagine production

growth continuing to occur into the future, with implications for both global hydrocarbon supply and a potential clean energy transition.

Measuring drilling productivity gains conditional on site quality is challenging because the only way to truly know a location's quality is to drill, and observed drilling locations are chosen endogenously. Firms hold rich but unobserved beliefs about local geology and learn over time from the outcomes of their own wells. In a standard Hotelling (1931) model, firms drill the highest quality locations first (Herfindahl, 1967). And in a model of resource exploration, firms learn from initial drilling and converge over time to the highest quality locations (Agerton, 2020; Hodgson, 2024). A naïve regression of well output on time will therefore conflate true productivity change with dynamic selection on site quality and learning: if the Hotelling forces dominate, then true firm productivity growth could be even larger than observed growth in output per well, since firms are getting more output out of worse locations. On the other hand, if initial signals were relatively uninformative and learning dominates, then observed growth in output per well could be entirely driven by site selection, rather than productivity improvements or other changes in firms' operational decisions.

To estimate productivity growth conditional on site quality, we develop a joint econometric model of well-level production and drilling decisions. The model is designed around distinct strategies for addressing two forms of selection: firms' decision of when to drill the first well in a given area, and their decision to drill subsequent wells based on what they learned from the previous wells. We estimate our model using data on drilling, production, and leasing from two major U.S. shale plays: the Bakken in North Dakota and the Haynesville in Louisiana. These plays are conducive to our analysis because of the large number of wells drilled to date, and because mineral leasing and drilling rights are organized into uniform administratively defined shapes, which allows us to characterize the universe of potential drilling opportunities from which the observed wells are selected. In addition, the Bakken primarily produces oil, while the Haynesville primarily produces natural gas, so that the two plays are exposed to different price variation over time and thus exhibit different boom-bust cycles of drilling during our sample (see figure 3).

We develop our model starting with just the first well drilled in each potential location. Output is a function of both site quality and a flexible function of time that captures the evolution of firms' operational decisions and technology adoption. A concern with estimating the output equation alone is that firms, with millions of dollars at stake for each well, may possess unobserved (to the econometrician) information about geologic quality for a specific drilling location, and decide whether and when to drill that location based in part on that information. We address this possibility through a joint model of output and selection into whether and when each potential location is drilled.

A key institutional feature that we leverage for identification is that before firms are able to drill, they must first acquire leases for the land on which they will drill. The time and expense of this acquisition process can vary considerably across proximate locations, based on the extent of parcel fragmentation and on mineral owners' willingness to lease (Leonard and Parker, 2020). This variation in leasing difficulty, and ultimately the timing with which leases are signed, creates variation in the timing of drilling across locations, even after controlling for observable measures of geologic quality.

To use this variation, we estimate a hazard model for drilling in which the timing of when each square-mile "section" of land is drilled is a function of when the section's leases are signed, along with an observable measure of geologic quality and calendar time. We then form a control function from the residuals of the survival function that is implied by the hazard. Including this control function in the output equation allows us to purge it of effects induced by site selection and thereby isolate the evolution of output growth over time, conditional on site quality. This approach is similar in spirit to the control function strategies discussed in Imbens and Newey (2009) and Newey (2009), but tailored here to address firms' selection of *when* to drill, not just whether to drill, at a particular location. We are able to use this control function approach rather than specify a full model of firms' optimizing behavior, such as the optimal stopping problems studied in Bhattacharya, Ordin and Roberts (2018), Agerton (2020), Herrnstadt, Kellogg and Lewis (2024), and Ordin (2024), because we are interested in understanding the evolution of output growth over time rather than in simulating counterfactuals that would alter firms' drilling decisions. This modeling choice also economizes on computational complexity and the strength of the assumptions that we need to impose.

The control function that we estimate is consistent with positive selection on unobserved quality. That is, unobsersably productive parcels are drilled earlier. As a result, we estimate larger productivity growth among first wells in each state, relative to estimates from a regression model without the selection correction, since locations drilled earlier in the sample are higher quality than later locations.

By the mid 2010's, most new wells in the Bakken and the Haynesville come not from expanding into new sections, but from putting additional wells in already drilled sections. As discussed in Agerton (2020), if firms initially have only a noisy signal of the section's quality prior to drilling, they will likely update their beliefs after observing production outcomes, and those beliefs can then influence their decision of whether to drill an additional well in the same section. We find evidence consistent with this form of learning in the data: firms are more likely to drill a second well in a section if the first well's output is higher than expected.

This learning mechanism invalidates a potential empirical strategy that is tempting by virtue of its simplicity. Suppose that we attempted to estimate productivity growth by using a model of well-level output that controlled for spatial variation in geology using granular section fixed effects, leveraging the fact that we see multiple wells drilled into the same section at different points in time. Such an approach has appeal because it obviates the need to address firms' private information before drilling the first well in each section. However, because each well's output contains both signal and noise, the outputs of drilling investments made using a cutoff rule on the first well's production will exhibit mean reversion. This mean reversion will then cause the simple fixed effects regression to produce an underestimate of true productivity growth.[1]

We therefore account for firms' learning by specifying a Bayesian updating model in which expected output and the firm's drilling decision are influenced by previously drilled wells' production shocks, relative to expectation. This model amounts to including a second control function in the outcome equation that accounts for the possibility that the production shock for the jth well in a given section will be correlated with the shocks from wells 1 through $j-1$.[2] Unlike Agerton (2020), we allow learning to be incomplete, motivated by the fact that well-to-well production outcomes are highly noisy. The identifying assumption in this model (and in Agerton (2020)) is that after the first well is drilled, the only factors that jointly affect subsequent wells' drilling decisions and the output from those wells are the outcomes of previously drilled wells. We estimate this learning model jointly with the first-well outcome equation (including the leasing-based control function), using GMM.

The full model allows us to decompose changes over time in average output per well along several margins. We find that unobserved site quality in both states declined over time, reducing wells' output approximately 0.1 log points. Moving in the other direction, learning from the outcomes of previously drilled wells led to firms to select sites that were 0.04 log points better post-2012. This effect from learning is an order of magnitude smaller than that found for the Haynesville in Agerton (2020), which assumed that firms perfectly learn about geologic quality after drilling each section's first well and then optimally choose locations for subsequent wells based on this information.[3] Finally, we find that later wells are 0.05 to 0.12 log points less productive due to depletion and well interference effects stemming

[1]This mechanism for bias in a simple section fixed effect regression is similar to the mechanism described in Lazear (2004) to explain declines in employee performance following promotions.

[2]These shocks are defined relative to the firm's expectation of output before the first well was drilled. We define the shocks precisely in section 3.2.

[3]Unlike our decompostion of output growth, which takes firms' drilling decisions as data, Agerton (2020) decomposes output growth into productivity versus learning by simulating firms' drilling decisions (taking leases and the time series of natural gas prices as data), conditional on their ability to perfectly learn geology after they drill the first well in each section.

from placing more and more wells within a single section. While each of these forces is economically meaningful, their net effect explains only a small fraction of the increased output growth over time. After accounting for all these forms of site selection, we find that firms' operational decisions, broadly defined, caused output per well to grow by roughly one log point in North Dakota between 2010 and 2019, and by more than that in Louisiana, conditional on site quality.

We next unpack the potential mechanisms behind this large increase in output that is unaccounted for by changes in site selection over time. We first demonstrate that it is driven by output improvements within existing firms rather than reallocation of drilling sites to high-productivity firms. Second, we show that it is associated with large increases in a subset of drilling and fracking inputs that we are able to observe: wells' horizontal lengths, and the volumes of water and sand used during hydraulic fracturing. We stop short of estimating causal impacts of these observable inputs on well-level output, however, since drilling and completing shale wells involves a multitude of unobserved inputs, such as expenditures on surface and downhole equipment, that are costly and likely complementary with the observed inputs. Instead, we study data on input prices and the total cost of wells' drilling and completion in order to discern whether oil and gas output increased simply because firms used more inputs, holding productivity constant, or whether there were process improvements or adoption of technologies that improved productivity and were complementary to the observed inputs. We find that the increases in input use cannot be explained by changes in input prices: while prices for inputs fell during our sample period, these changes were smaller in magnitude than contemporaneous large decreaes in prices for oil and natural gas. Reported all-in drilling and completion costs track input price changes but with a dampened magnitude, and they exhibit no evidence of a large increase commensurate with our estimated improvements in output per well. Thus, the evidence implies that firms substantially increased their well-level input use and oil and gas production, conditional on site quality, while holding their total costs relatively flat.

We therefore conclude that the large increase in well-level oil and gas production during the shale boom stemmed primarily from productivity improvements, and that these improvements were complementary to increases in well-level input use. While we also find evidence of changes over time in the average quality of the sites drilled, these effects on net push against output growth, rather than in favor of it. These selection effects are not trivial and amount to roughly a 0.1 log point decrease in output per well over the sample period — a magnitude sufficient to be a major force in industries that exhibit "typical" rates of productivity improvement on the order of 1–2% per year. But in shale oil and gas, these site selection effects are swamped by the large productivity gains and associated operational

decisions that drove output growth during the boom.

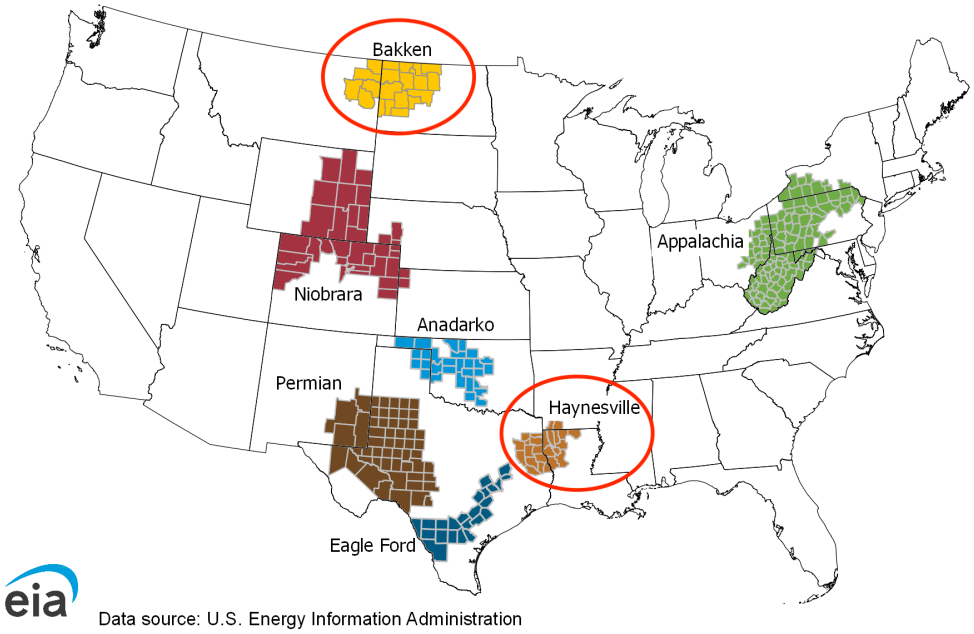
The remainder of the paper proceeds as follows. Section 2 begins by discussing relevant institutional details of the shale industry and highlighting descriptive features of our data. Section 3 presents our joint model of output and firms' drilling decisions, first emphasizing the model's identification and then turning to estimation. Section 4 then presents our main decomposition of changes over time in output per well into site selection versus technology and operational decisions, and section 5 follows by showing how data on input use, prices, and costs point to a story of productivity improvement. Finally, section 6 concludes by summarizing our results, suggesting a direction for future work, and discussing how our model might be applied to other settings in which the outcomes from firms' investments depend on both firms' ability to select high-quality investment targets and their ability to execute, conditional on the target.

2 Background and data

Shale oil and gas formations in the United States are typically located several thousand feet below the surface—often roughly 6,000 to 12,000 feet deep—similar to the depths of many conventional reservoirs. What distinguishes shale formations from conventional reservoirs is that in shale, the hydrocarbons are trapped in low-permeability rock. In conventional reservoirs, oil or gas accumulates in porous formations that allow hydrocarbons to flow easily toward a vertical well once the reservoir is drilled. In shale formations, by contrast, hydrocarbons are dispersed within rock with extremely low permeability, meaning that they do not flow freely to a wellbore. As a result, production from shale requires technologies that both expose a large surface area of the rock to the well and create pathways through which hydrocarbons can flow.

Modern shale development therefore relies on a combination of two technologies, horizontal drilling and hydraulic fracturing, that together enable economic extraction from formations that were previously considered uneconomic. Horizontal drilling allows operators to steer the well laterally through the target formation rather than drilling only vertically, greatly increasing the length of the well that is in contact with the hydrocarbon-bearing rock. While directional drilling technologies date back several decades, long horizontal laterals became widely deployed in U.S. shale plays during the late 1990s and early 2000s. Hydraulic fracturing ("fracking") involves injecting high-pressure mixtures of water, sand, and chemicals into the well to create fractures in the surrounding rock; the sand keeps these fractures open, allowing hydrocarbons to flow toward the well. Fracturing techniques were first developed in the late 1940s and used for decades in conventional wells, but the high-

Figure 2: Map of Major Shale Plays



Source: U.S. Energy Information Administration.

volume, multi-stage fracturing methods used in shale became commercially viable in the 2000s, when they were combined with horizontal drilling to unlock large shale resources.

Several technological and operational changes may have contributed to improvements in this process since the start of the shale boom. First, operators have substantially increased the intensity of hydraulic fracturing treatments. Modern shale wells use much larger volumes of water and proppant (sand) than early wells, which increases the size and conductivity of the fracture network and exposes more of the reservoir to the wellbore (Covert, 2015; Steck, 2022). Second, horizontal laterals have become longer, allowing each well to contact a larger portion of the formation and increasing the amount of reservoir drained by a single well. Third, firms have improved the placement and design of fracture stages along the lateral, including better targeting of geologic "sweet spots" within the formation and more sophisticated multi-stage fracturing techniques that create a denser fracture network (Curtis, 2016; Casero, 2021). More recently, operators have also begun using machine learning and advanced data analytics to optimize completion design control drilling operations (Wethe, 2024). In addition, productivity gains may arise from improved drilling efficiency, better well spacing and pad drilling practices, and learning-by-doing as firms accumulate experience operating within specific shale plays.

2.1 Well Data

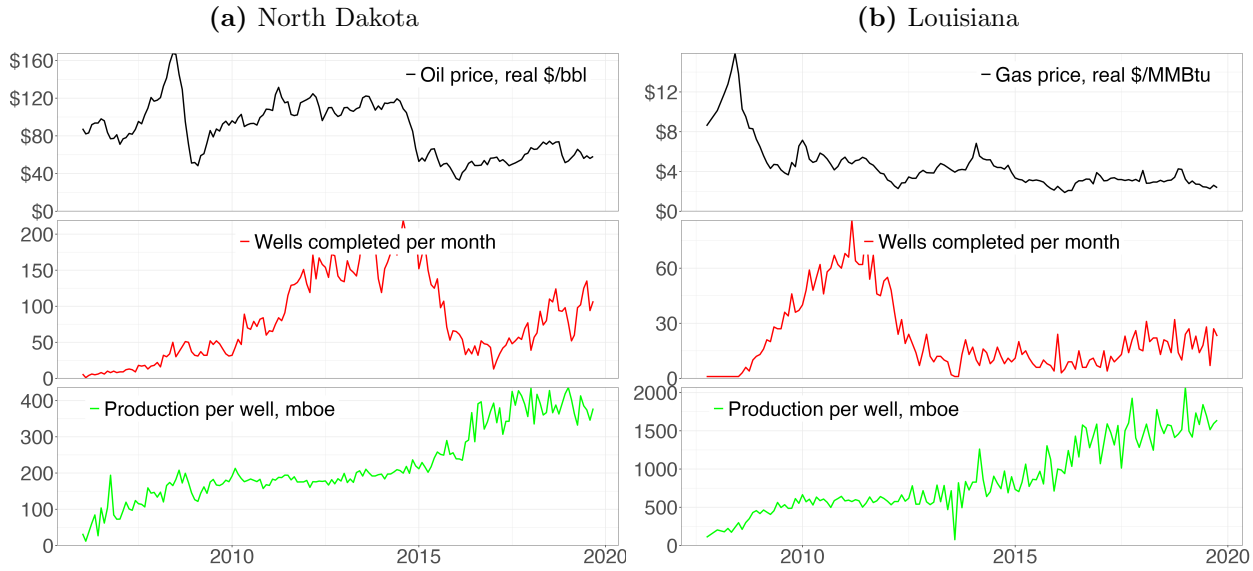
We use well data from Enverus, a commercial data vendor that compiles information on drilling, production, and leasing from state regulatory agencies. For each well, we observe its drilling date, completion date, location (latitude and longitude), vertical depth and horizontal length. We also observe the original firm ("operator") that drilled the well. These data include both modern shale wells and conventional wells. To restrict our analysis to shale wells, we use the well's depth and horizontal length, as well as string matching on the well's target formation. We focus our attention on two of the U.S.'s major shale plays: the Bakken in North Dakota and the Haynesville in Louisiana (figure 2).[4] The Bakken primarily produces oil, while the Haynesville primarily produces gas.

We then merge in production data for each well, also obtained from Enverus. For every well, we observe the monthly volumes of gas and oil produced, in units of millions of British thermal units (mmBtu) and barrels (bbl), respectively. We combine these volumes into a single barrels of oil equivalent (boe) measure by dividing reported gas output by 5.67, the energy content of crude oil in mmBtu per bbl. We then aggregate this monthly production data into cumulative measures of production in the first 12 and 36 months. These truncated measures are useful because our production dataset ends in 2023, and because a common rule of thumb in the industry is that modern shale wells will produce approximately 70 percent of their lifetime production in the first 36 months. Moreover, lifetime production is typically assumed to follow a hyperbolic Arps "decline curve" in which the ultimate oil recovered is proportional to output in the first 36 months. We limit our well sample to wells drilled by the end of 2019, so that we have at least three years of production data for all wells in our sample.

Figure 3 plots, for each month in the sample, the number of new wells completed and the average 36-month production of new wells, for both the Bakken (North Dakota) and the Haynesville (Louisiana). The top panels plot the average oil and gas spot prices. Both price series drop precipitously during the great recession. The oil price recovers shortly thereafter but collapses in late 2014, primarily driven by a decrease in global oil demand (Baumeister and Kilian, 2016). Natural gas prices fluctuate but overall decrease throughout the sample period. The number of new wells drilled declines noticeably in 2012 in the Haynesville, while the Bakken continues to see growth in drilling until oil prices fall in late 2014. The average production of new wells increases substantially over time in both plays.

[4]Our initial focus on these plays is primarily driven by data and institutional considerations. In both, spacing units are typically defined as Public Land Survey System (PLSS) sections, which facilitates our ability to use leasing timing in our research design. In contrast, spacing units are irregularly defined in Pennsylvania and Texas.

Figure 3: Price, Completions, and Production Over Time



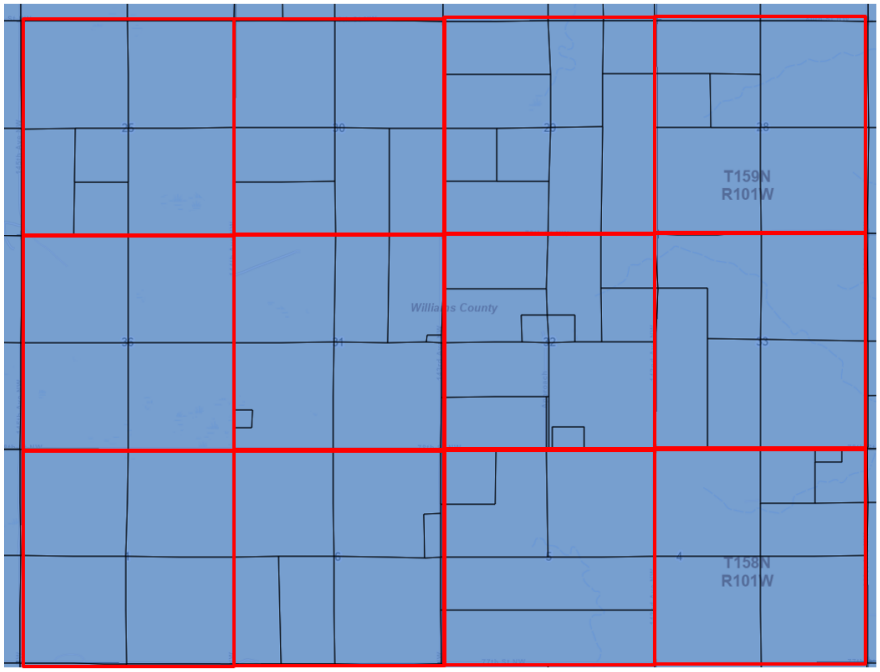
Notes: The top panel shows average real oil prices (West Texas Intermediate) and average real natural gas prices (Henry Hub) over time. The middle panel shows completions per month in North Dakota and Louisiana. The bottom panel shows total monthly production for new wells by month. See section 2 for details on data sources and construction.

2.2 Land and lease data

Mineral rights in the onshore United States are primarily owned by private individuals (Covert and Sweeney, 2023; Herrnstadt et al., 2024). In order to monetize these resources, mineral owners form partnerships with oil exploration and production companies ("E&Ps") through lease agreements. These leases give E&Ps the option, but not the obligation, to drill and produce on the land. Leases contain three core contractual elements: an up front bonus payment, a royalty rate specifying the share of revenue paid to the mineral owner, and a "primary term" that specifies the date by which drilling must commence in order for the lease to remain active. Primary terms are typically three to five years, and leases will sometimes also include an option for the lessee to extend the lease for another two years in exchange for a flat payment. If a producing well is drilled before the end of the term, then the E&P typically retains the right to drill additional wells indefinitely. Lease timing thus imposes hard timing constraints on the timing of drilling, as it cannot occur before a firm has secured drilling rights, and it must occur before the end of the lease term.

In both North Dakota and Louisiana, oil and gas development is organized around the Public Land Survey System (PLSS), which divides land into a regular grid of "sections." The PLSS was established by the federal government in the late eighteenth century to facilitate the sale and settlement of public lands in the United States. Under this system,

Figure 4: Example Section and Parcel Map



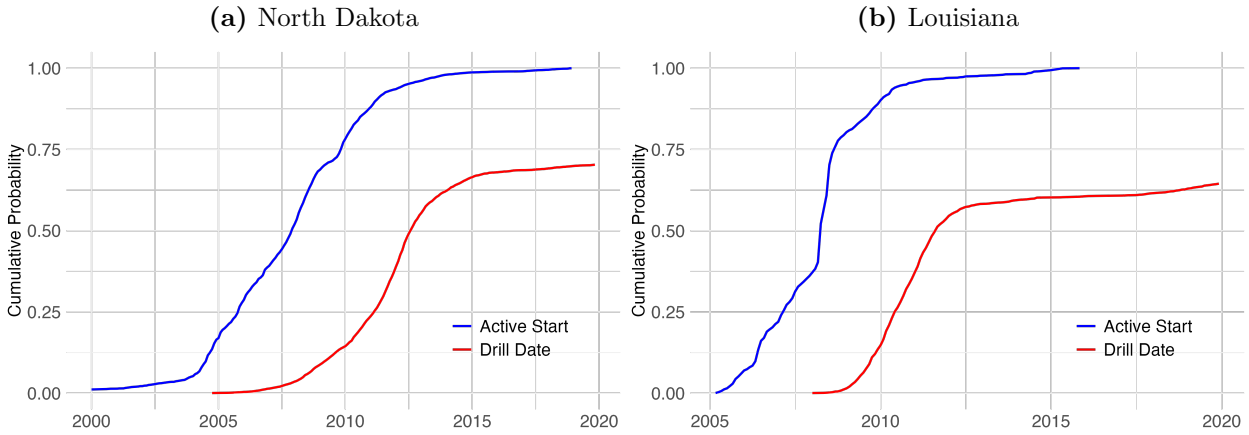
Notes: Parcel map from the North Dakota GISHUB. The red boxes indicate sections, and the black lines indicate the boundaries of individual parcels that are owned by different mineral owners.

land is divided into townships that are six miles by six miles, each of which contains 36 sections arranged in a one-mile by one-mile grid. Each section therefore contains roughly 640 acres, and can be further subdivided into quarter-sections (160 acres) and smaller parcels for ownership and leasing purposes.

Sections play an important role in both mineral leasing and drilling decisions. Because mineral ownership is often fragmented across many landowners within a section, companies must assemble leases covering a sufficient share of the acreage before drilling can occur. Regulatory authorities in both states generally define drilling or "spacing" units at the section level (or multiples of sections), which determines the area over which production is pooled and royalties are distributed. As a result, the section serves as a natural geographic unit for organizing leases, determining drilling locations, and allocating production across mineral owners. Figure 4 shows an example of a typical land pattern in the Bakken, wherein sections form a regular grid, and parcels have varying shapes and sizes within each section (though quarter-section parcels are fairly common, at least outside of towns). One feature of this map that will later be useful for our research design is that the degree of parcel fragmentation varies substantially across adjacent sections.

From Enverus, we obtained data on all leases in both North Dakota and Louisiana. For

Figure 5: Section lease and drill timing



Notes: Figure shows the cumulative distribution of the date a section if 50% leased and the date of the first well drilled.

each lease, we observe the lessor, lessee, term dates and royalty rate. We do not observe bonus payments. For North Dakota, we use shapefiles for each lease to map leases to sections spatially. For Louisiana, we are able to map leases to specific sections using township-range-section indicators in the tabular lease data. For both plays, we drop duplicates based on having identical locations, lessees, parcel areas, and term dates. Even after this de-duplication, there remain many sections in which the total leased area exceeds the section area (640 acres) for a meaningful interval of time. While this excess may reflect in part some mis-measurement in the lease data, we believe that much of the excess reflects parcels for which mineral ownership is held in multiple "undivided" interests. That is, for many parcels, multiple owners (typically progeny of the original owner) hold fractional ownership of the entire parcel, and in these cases, the lessee must contract with each of them to fully secure drilling rights on the parcel. Each such contract appears as a separate record in our lease data, thereby leading to situations in which leased acreage can (sometimes greatly) exceed section acreage.

2.3 Sample construction

Our unit of observation for modeling investment is a section. We first restrict the sample to sections in North Dakota and Louisiana that are located within the boundaries of the Bakken and Haynesville shale plays (as defined by the Energy Information Administration in figure 2). In our primary analysis, we further restrict the sample to sections with the convex hull of wells in this subset, thereby removing areas that are not part of the main development regions. Next we compute the running share of each section under lease in each month. We

Table 1: Sample Summary Statistics

(a) North Dakota

	Mean	P10	Median	P90
OOIP	7.05	4.0	7.0	11.00
Wells per section	3.76	1.0	2.0	9.00
Mean BOE (36 mo)	183.65	71.8	158.8	328.54

(b) Louisiana

	Mean	P10	Median	P90
OGIP	101.52	47.11	102.90	147.54
Wells per section	2.75	1.00	2.00	6.00
Mean BOE (36 mo)	772.34	285.55	667.46	1426.62

Notes: Summary statistics for the final section sample used in the analysis. In each panel, the unit of observation in the first row is a section, the unit of observation in the second row is a section with at least one completed well, and the unit of observation in the third row is a completed well. The units for OOIP are millions of bbl per square mile, and the units for OGIP are billions of cubic feet per square mile. OOIP is recorded as integer data. See section 2 for details on data sources and construction.

define a section as "active" in a given month if at least 50 percent of its acreage is under lease, and we restrict attention to sections that become active between 2000 and 2019.[5] The resulting sample represents the potential drilling opportunities available to firms in the sense that the sections therein: (1) overlay productive regions of known shale formations; and (2) both landowners and firms have taken the initial step of signing mineral leases.

We use geospatial information on the path of each well to spatially match wells with sections.[6] For each section, we define the date it is drilled as the date for which a completed well first commences production in the section. Figure 5 plots the cumulative distribution function (CDF) of the date by which sections in each state become "active" (50% leased). In both states leasing is minimal prior to 2005 and then largely complete by 2010. The red lines plot the CDF of the date that the first well in each section is drilled. Consistent with the timing of leasing and the limits of primary terms, the bulk of drilling occurs three to five

[5]For sections in which the total leased area is always below the section's area, we define the section as active when its total leased area first reaches 50% of the section area. For sections in which the total leased area exceeds, at any time, the section's area, we define the section as active when its total leased area first reaches 50% of the maximum leased area.

[6]Many well laterals intersect two or more sections; each such well will account for multiple observations in our final estimation sample. All of our reported standard errors are clustered on township, which will account for the high correlation between the residuals for these observations.

years after leasing. Roughly 30% of options to drill in sections go unexercised by the end of our sample, and there are very few sections for which a first well is drilled later than 2015. Table 1 presents summary statistics on sample sections in each state. Conditional on any drilling, the average section has 3.75 wells in the Bakken and 2.75 wells in the Haynesville. In both plays, the distribution of wells per section is substantially right-skewed.

For each section, we have collected a measure of geologic quality from the literature on petroleum geology. For the Haynesville shale in Louisiana, we follow Agerton (2020) and Herrnstadt et al. (2024) by using estimates of "original gas in place" (OGIP) from Gülen, Ikonnikova, Browning, Smye and Tinker (2015), and for the North Dakota Bakken, we use estimates of "original oil in place" (OOIP) from Amin Gherabati, Browning, Male, Hamlin, Smye, Walsh, Ikonnikova, McDaid and Lemons (2017). These estimates are based on data on how formation thickness, porosity, temperature, and pressure vary across space. For the Bakken, we were unable to obtain the section-by-section estimates from the authors, so we instead digitized the OOIP maps shown in figures 17 and 18 of Amin Gherabati et al. (2017). As a consequence, our OOIP measure for the Bakken is an ordered categorical variable that runs from 1 to 14.[7] Summary statistics for these quality measures are provided in the bottom row of each panel of table 1.

Finally, we describe the distribution of firms operating in our sample. Table 2 tabulates the share of wells drilled by the top four firms in each state. In North Dakota, each of the top three firms drilled around 10% of wells in the state. In Louisiana, the top firm (Chesapeake) drilled more than 30% of wells during the sample period. Figure 6 plots the within-year Herfindahl-Hirschman Index (HHI) to evaluate how the market concentration of drilling changed over time. In North Dakota, the HHI is low—almost always below 0.1—and stable throughout the sample, indicating that the industry was not concentrated. In Louisiana, the HHI is higher and more volatile, with a noticeable spike around 2015, suggesting some potential scope for reallocation of drilling toward among firms during our sample period.

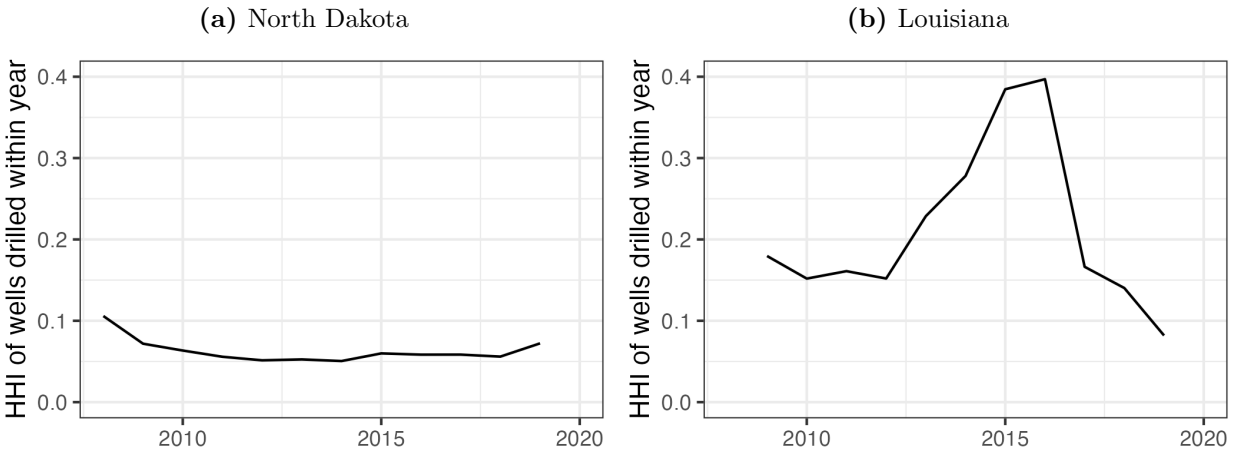
This descriptive evidence suggests that changes in the allocation of wells across firms are unlikely to have been important to the evolution of productivity in North Dakota, but may have been in Louisiana. Given the overall low market shares of all firms other than Chesapeake in Lousiana, our primary model abstracts away from firm-level heterogeneity in both states. We explore the potential role of allocation of drilling opportunities to specific firms later in section 5.1.

[7]The Bakken shale consists of two main producing formations that are stacked vertically: the Middle Bakken and the Three Forks. Figure 17 in Amin Gherabati et al. (2017) maps the OOIP for the Middle Bakken, with colors indicating when OOIP is $<$ 1,1-2,...,5-6,$>$ 6 million bbl per square mile. Figure 18 does the same for the Three Forks. We digitized these maps, and for each section our OOIP variable is the sum of the OOIP categories from the two maps.

Table 2: Share of Wells by Top 4 Firms

North Dakota		Louisiana	
Firm	Share	Firm	Share
Continental Resources	12.4%	Chesapeake	31.6%
Hess	9.7%	BHP Billiton	15.5%
Whiting	8.7%	Encana Oil & Gas	10.9%
EOG Resources	4.0%	SWEPI	5.1%

Figure 6: Herfindahl-Hirschman Index (HHI) of Drilling Firms over Time



Notes: The plotted range for North Dakota is from 2008, and for Louisiana, from 2009, reflecting that most wells are drilled after these years.

3 Model and estimation

This section presents the econometric model that we use to identify and estimate the contributions of site selection to growth in output per well over time. Sub-section 3.1 begins by discussing our model of selection into whether and when a drilling opportunity (henceforth, a "section") receives its first well. Sub-section 3.2 then discusses our model that accounts for how firms' drilling decisions may be affected by their learning about geological quality through observing production outcomes from earlier wells. Sub-section 3.3 then discusses how we estimate the model.

3.1 First-well selection and timing

For each section i, we model y_{i1}, the log barrels of oil equivalent (boe) produced by the first well $j = 1$ as

$$y_{i1} \;=\; \omega(t_1(i)) + X_i'\beta + \eta_i + \theta_i + \varepsilon_{i1}, \tag{1}$$

where $t_1(i)$ is the first well's drilling date. X_i includes observed geologic characteristics that we will operationalize with our OOIP and OGIP quality measures. η_i denotes a scalar section-level productivity shifter that captures the firm's unbiased belief about the section's quality prior to drilling. η_i is unobserved by the econometrician and by construction is mean zero conditional on X_i. θ_i is the difference between section i's true quality and the firm's ex ante belief η_i, and by construction is mean zero conditional on X_i and η_i. Finally, ε_{i1} is a mean-zero well-specific idiosyncratic output shock that is realized after the well is completed. Neither θ_i nor ε_{i1} can affect the firm's first-well drilling decision, and these terms could be collapsed into a scalar composite error for the purpose of this sub-section. Later in sub-section 3.2 we will allow the firm to learn about θ_i by observing the output of the first and subsequent wells.

We are interested in the evolution of ω, which reflects operational choices and technological advances that increase output per well conditional on site quality. While we could in principle estimate time-period fixed effects, in practice we model $\omega(t)$ a flexible function of calendar time. If the distribution of η_i is non-degenerate, then $t_1(i)$ will in general be correlated with η_i, as will the selection of section i into being drilled at all (i.e. whether $t_1(i) \leq T$, where T is the last period of observation). Consequently, a simple OLS estimate of the outcome equation (1) will not recover the true time path $\omega(t)$.

To address selection into whether and when each section i is drilled, we model the drilling timing for i's first well:

$$t_1(i) = g\big(X_i, Z_i, u_i\big), \tag{2}$$

where Z_i summarizes the difficulty of assembling leases in section i, and u_i is a scalar unobserved component that is potentially correlated with η_i and is the source of endogeneity in the output equation (1).

Let F_u denote the CDF of u_i. We then define $V_i = F_{t_1|X,Z}(t_1(i)|X_i, Z_i) = F_u(u_i)$, so that V_i is the conditional probability of drilling the first well in section i by its observed drilling time $t_1(i)$, given X_i and Z_i. Intuitively, V_i captures how the variable u_i, which is known to the firm but not the econometrician, drives drilling timing through the function g. Conditioning on it in a control function will address the endogeneity of $t_1(i)$ and allow us to

identify $\omega(t)$.

The control variable approach that we adopt, following Imbens and Newey (2009), makes the following assumptions:

1. (η_i, u_i) and (X_i, Z_i) are independent.
2. u_i is continuously distributed, F_u is strictly increasing, and g is strictly monotonic in u_i.
3. For all (X_i, t_i), the support of V_i conditional on (X_i, t_i) equals the support of V_i.

Per Imbens and Newey (2009), the first two assumptions are sufficient for the result that $t_1(i)$ and η_i are independent conditional on V_i and X_i. Intuitively, conditioning on V_i is equivalent to conditioning on u_i, since V_i is a one-to-one function of u_i. Then, once u_i is conditioned on, along with the conditioning on X_i, all remaining variation in $t_1(i)$ must be driven by variation in the instrument Z_i. The third assumption is analagous to (though stronger than) a rank condition from a linear simultaneous equations model. It requires sufficient variation in Z_i to shift (the distribution of) drilling timing across the entire calendar time window of interest for any given value of X_i.

We may then write the expected log output of the first well in each section i, conditional on X_i, $t_1(i)$, and V_i as:

$$E[y_{i1}|t_1(i), X_i, V_i] = \omega(t_1(i)) + X_i'\beta + E[\eta_i|V_i], \tag{3}$$

where the inclusion of the $E[\eta_i|V_i]$ term allows for consistent estimation of β and the function $\omega(t_1(i))$.

There are several ways to model the function $t_1(i) = g\big(X_i, Z_i, u_i\big)$. We choose an approach based on a hazard model. In continuous time, denote the instantaneous hazard for drilling section i at some time t as $h\big(t, X_i, Z_i\big)$. The hazard then implies the following survival model as the implicit function $t_1(i) = g\big(X_i, Z_i, u_i\big)$:

$$u_i = F_u^{-1}(V_i) = 1 - \exp\left(-\int_0^{t_1(i)} h\big(s, X_i, Z_i\big)ds\right). \tag{4}$$

Note that monotonicity of g with respect to u_i is ensured by the additive separability of u_i in survival model of expression (4).

An alternative approach would be to model the drilling timing decision as an optimal stopping problem and estimate the parameters of that problem jointly with the output equation, as done in Bhattacharya et al. (2018), Agerton (2020), Herrnstadt et al. (2024), and

Ordin (2024).[8] We eschew that approach here because: (1) our primary concern is measuring productivity, rather than evaluating counterfactuals that would alter firms' drilling decisions (as in the papers cited above); (2) solving and estimating the firm's optimal stopping problem would be computationally intensive; and (3) solving this problem would require us to specify the firm's beliefs regarding the evolution of productivity and other profit shifters (in addition to specifying the output function as in our current approach), and mis-specifying these beliefs might cause our productivity estimates to be inconsistent.

3.2 Subsequent-well selection, learning, and depletion effects

As shown in table 1, many sections in our sample are ultimately drilled more than once. This section augments the model developed above to account for two mechanisms that might affect the output from wells drilled after the first, beyond just the evolution of $\omega(t)$. First, as more and more wells are drilled within a square-mile section, the output of later wells may be affected by depletion and interference from previous wells. Second, firms' decisions about whether and when to drill subsequent wells in a section may be affected by their learning about the section's latent productivity shifter θ_i after drilling the first well. This addition to the model is important because, as we will see in section 4.2 below, the data show that the probability that a section receives a second well is strongly related to the "surprise" component of the first well's output, which we define as the residual from our estimate of equation (3) above. This pattern is consistent with behavior in which firms learn over time which sections are truly high-quality, and then drill those sections relatively frequently. If we were to ignore this process and instead simply estimate equation (3) for all wells rather than just the first well, we would overstate the extent to which increases in output per well over time are attributable to the mechanisms underpinning $\omega(t)$.

In our multiple-well model, the output of the jth well in section i is given by

$$y_{ij} \;=\; \omega(t_j(i)) + X_i'\beta + \lambda_j + \eta_i + \theta_i + \varepsilon_{ij}, \tag{5}$$

where $t_j(i)$ is the drilling date of the jth well, and the new λ_j term is a jth well fixed effect that accounts for depletion and well interference ($\lambda_1 = 0$ by construction).

To see how an estimate of equation (5) will be biased if learning is not accounted for, consider the expected output of the second well $j = 2$, which is only observed if a second well

[8] Yet another alternative would be to estimate the output equation (1) by 2SLS, using Z_i as an instrument for $t_1(i)$. One drawback of that approach is that it would not account for selection into whether a given section is drilled at all during our sample period. Another is that it would not allow us to flexibly identify $\omega(t)$, which the control function approach accomplishes by taking advantage of the fact that the endogeneity problem is driven by a single index, η_i.

is drilled. Define e_{i1} as the first well's production shock: $e_{i1} = y_{i1} - \omega(t_1(i)) - X_i'\beta - E[\eta_i|V_i]$. We then have the expectation:

$$E[y_{i2}|t_2(i), X_i, \eta_i, y_{i1}] = \omega(t_2(i)) + X_i'\beta + \lambda_2 + E[\eta_i|V_i] + E[\theta_i|e_{i1}].$$

If η_i is only an imperfect signal of the true underlying geology, then the difference θ_i between the section's true quality and η_i will be correlated with the shock e_{i1}, so that $E[\theta_i|e_{i1}]$ is a strictly increasing function of e_{i1}. And if the firm is aware of the imperfection of its initial signal η_i and learns from the first well's output, then a large realization of e_{i1} will make it more likely that the second well is drilled, and accelerate its timing conditional on being drilled. Because second (and subsequent) wells are drilled after first wells by construction, this pattern will lead a least-squares estimate of $\omega(t)$ from equation (5) to overstate the increase in productivity over time.[9]

Extending the model to the jth well, let e_i^{j-1} denote the history of output shocks (relative to the firm's ex ante expecation before drilling the first well but accounting for the depletion effects λ_j) for all previous wells in the section: $e_i^{j-1} = \{e_{i1}, \ldots, e_{i,j-1}\}$. The expected output of the jth well, conditional on the history of output shocks from the previous $j-1$ wells, is

$$E[y_{ij}|t_j(i), X_i, \eta_i, y_i^{j-1}] = \omega(t_j(i)) + X_i'\beta + \lambda_j + E[\eta_i|V_i] + E[\theta_i|e_i^{j-1}]. \tag{6}$$

To consistently estimate equation (6), we address learning by conditioning on previous wells' output shocks, thereby accounting for the $E[e_{ij}|e_i^{j-1}]$ term. The identifying assumption is that after drilling the first well, firms only update their beliefs about section i's latent productivity based on the output from the previous wells. Agerton's (2020) study of productivity and learning in the Haynesville shale also imposes this assumption, along with the stronger assumption that firms fully learn the true quality θ_i after drilling the first well. Our model relaxes this assumption by allowing for incomplete learning, motivated by the substantial well-to-well noise in output that we observe. Consistent with the idea that each well's output is only a noisy signal, we shall see in the results section 4.2 below that the second well output shocks e_{i2} are positively correlated with the first well shocks e_{i1} but strongly mean-reverting.

[9]If the $E[\theta_i|e_{i1}]$ term is not accounted for, the well order fixed effects λ_j will capture some of the learning effect, causing the λ_j to increase with j relative to the case when these terms capture depletion effects alone. The estimates of the terms in the output equation (including $\omega(t)$) will still be biased, however, because the effects of learning on expected output and drilling timing are a function of the realized output shock from the first well, not just whether or not a subsequent well was drilled.

3.3 Estimation

For the instrument Z_i, we use the date by which at least half of the acreage in section i is under lease (denoted τ_i). Sections with many distinct mineral owners or more reluctant lessors take longer to reach this threshold, so τ_i shifts the timing of the first well. To address the possibility that τ_i may partially reflect variation in geologic quality, we residualize τ_i on fine grid fixed effects in latitude and longitude, so that the identifying variation in τ_i comes from local differences in leasing difficulty rather than from broad geologic trends.

We estimate the hazard as a piecewise-exponential proportional hazards model on monthly-level section drilling data, using a Poisson GLM with a flexible baseline hazard captured by natural cubic splines in t. This Poisson representation is the standard likelihood-equivalent implementation of piecewise-exponential proportional hazards models. The fitted interval-specific hazards are treated as piecewise constant within a month, and we recover the implied CDF of t by exponentiating the cumulative hazard. In the output equation, we use the same natural cubic spline basis for t as in the hazard model to recover $\omega(t)$, controlling for a polynomial function H of the estimated CDF at the realized drilling date, $H(\hat{V}_i) = H(\hat{F}\big(t_1(i)|X_i, Z_i\big))$, as an additional regressor to correct for selection.

To tractably estimate the learning model, we assume that firms follow Bayesian updating. Before the firm drills the first well in section i, let its prior about the section's true quality be $\theta_i \sim \mathcal{N}(0, \sigma_p^2)$. Let σ_s^2 denote the variance of the iid ex post output shocks ε_{ij}. Then, the posterior mean of θ_i before drilling the jth well is:

$$E[\theta_i|e_i^{j-1}] = \frac{j-1}{(j-1)+\sigma_s^2/\sigma_p^2} \cdot \bar{e}_i^{j-1}, \text{ where } \bar{e}_i^{j-1} = \frac{1}{j-1}\sum_{k=1}^{j-1} e_{ik}. \tag{7}$$

The formula clarifies that the signal noise relative to the prior noise, σ_s^2/σ_p^2, is identified, but not the levels of the signal and prior noise. Thus, we parameterize the learning process using a single parameter $\kappa = \sigma_s^2/\sigma_p^2$.[10]

Let $\hat{e}_{ij}$ denote the estimated output equation residuals (taken with respect to expected output prior to observing the first well's outcome), and let $\hat{\bar{e}}_i^{j-1}$ denote the average residual from all previous wells in the section. We then have the following estimation equation for

[10] We can allow for a more general learning process (e.g., non-normal distributions or non-Bayesian learning) and still keep the estimation tractable by assuming that $\bar{e}_i^{j-1}$ is a sufficient statistic for θ_i before drilling the jth well. Then, $E[\theta_i|e_i^{j-1}] = E[\theta_i|\bar{e}_i^{j-1}] = l_{j-1}(\bar{e}_i^{j-1})$, where l_{j-1} is a flexible but strictly increasing function, for any reasonable learning process, which can be estimated with a shape restriction for monotonicity for each j.

the jth well in section i:

$$y_{ij} = \omega(t_j(i)) + X_i'\beta + H(\hat{V}_i) + \lambda_j + \frac{j-1}{(j-1)+\kappa} \cdot \hat{\bar{e}}_i^{j-1} + \varepsilon_{ij}. \tag{8}$$

Equation (8) is a nonlinear in its parameters, both due to the functional form for κ and because $\hat{\bar{e}}_i^{j-1}$ is itself a function of all parameters. We therefore estimate it by GMM, using estimates from first wells as starting values. In practice, firms drill "batches" of wells contemporaneously within a section. We incorporate this feature into our estimation by treating wells in a batch as sharing the same posterior mean.[11] We also parameterize the well order fixed effects λ_j as a polynomial function of j.

4 Results

4.1 Estimates using only the first well in each section

We begin by estimating the output equation using only the first well in each section, accounting for selection into whether and when a section is drilled using the leasing-based control function. These estimates will then provide a foundation for estimating the full model that uses all wells and accounts for firms' learning from previous wells' output.

Our model requires a variable Z_i that shifts the timing of drilling but is independent of the unobserved section-level productivity shifter η_i, conditional on observed quality X_i. A natural candidate for Z_i, following Leonard and Parker (2020), is the extent to which parcel ownership was fragmented in section i prior to the shale boom. Fragmentation is plausibly exogenous, and the need to negotiate with many mineral owners will delay leasing, which in turn will delay drilling. We have experimented with data on parcel ownership from North Dakota's State Parcel Program,[12] but we found that these data are under-powered for predicting leasing because they only list at most two owners per parcel, while we know from our lease data that many parcels have far more than two owners.[13] Moreover, parcel ownership fragmentation alone will not capture variation across mineral owners in their willingness to lease.

[11] In our empirical implementation, wells in a section are grouped into a single batch if they are drilled within 60 days of the preceding wells. We impose this 60-day threshold because it is too short a period for a firm to observe a prior well's production outcome and incorporate that learning into the drilling operations of a subsequent well.

[12] North Dakota State Parcel Program data are available at https://gishubdata-ndgov.hub.arcgis.com/maps/ac6da1176038457db16e8debe3f1abaf/about.

[13] These situations take the form of "undivided mineral interests", wherein ownership in a parcel is shared among multiple owners, often members of the same family.

We therefore use the actual timing of leasing as our instrument, which captures both the effects of ownership fragmentation and of lessors' willingness to lease. Let τ_i denote the date by which at least half of the acreage in section i is under lease. τ_i exhibits substantial variation across sections in our sample regions. In North Dakota, τ_i spans from December 2003 at the 5th percentile to July 2012 at the 95th percentile, with a median of December 2007. Similarly, in Louisiana, τ_i ranges from November 2005 at the 5th percentile to August 2010 at the 95th percentile, with a median of April 2008. Furthermore, τ_i is a strong predictor of drilling timing in both states, yielding a correlation coefficient of 0.38 in North Dakota and 0.21 in Louisiana.

One might be concerned that τ_i reflects variation in unobserved quality η_i if firms more aggressively lease in sections that they believe are highly productive. While the correlation between τ_i and η_i is of course not testable, we can examine the correlation between τ_i and the observed quality measure X_i. In North Dakota, there is some correlation (correlation coefficient of 0.11), though the predictive power of X_i on τ_i is limited (leave-one-out out-of-sample R^2 of 0.013). In Louisiana, this relationship is even weaker, with effectively zero correlation between lease timing and observed quality (correlation coefficient of -0.02 and no predictive power (leave-one-out out-of-sample R^2 of -0.002).

Still, to better ensure that the identifying variation in τ_i comes from variation in leasing difficulty rather than variation in geology, we residualize τ_i on fixed effects for fine geographic grids that we overlay on our sample regions: a 20-by-20 grid for North Dakota and a 12-by-12 grid for Louisiana. This approach exploits the fact that there is considerable local variation in τ_i, whereas geologic variation is smoother over space. Indeed, the correlation between the residualized τ_i and X_i drops effectively to zero in both states: -0.003 (with a 95% confidence interval (CI) clustered by township of [-0.029, 0.023]) in North Dakota, and -0.010 (95% CI: [-0.085, 0.065]) in Louisiana. At the same time, the first-stage correlation with drilling timing remains positive and statistically significant, at 0.18 (95% CI: [0.15, 0.22]) in North Dakota and 0.13 (95% CI: [0.03, 0.23]) in Louisiana.

Table 3 presents the estimation results for our joint model of the first-well drilling timing (column "Hazard") and first-well output (column "First Well"). In the first-well drilling hazard model, the residualized lease start date exhibits a significant negative coefficient, confirming that localized leasing difficulties effectively delay drilling. In the first-well output model, the observed geologic quality variable positively affects production as expected. It is informative to visualize the shape of the control function implied by the selection model as in figure 7. The negative relationship between the predicted drilling probability (control variable $\hat{V}_i$) and the control function value indicates that sections that are drilled earlier are more likely to have higher unobserved quality, consistent with the idea that firms target the

most promising sites first and then move to less promising sites over time.

Table 3: Estimation Results

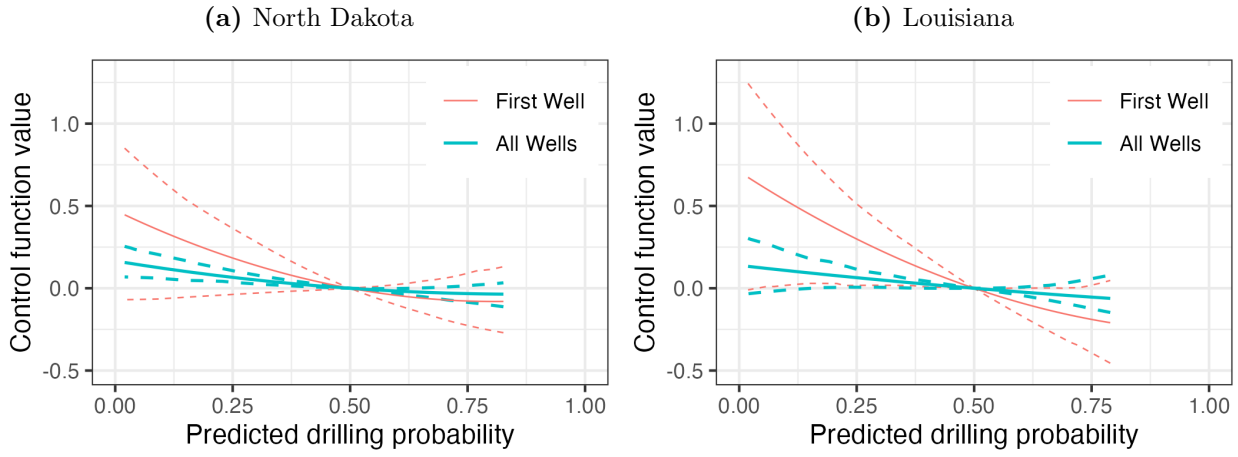
	North Dakota			Louisiana		
	Hazard	First Well	All Wells	Hazard	First Well	All Wells
Intercept	-9.539	9.590	9.387	-11.689	9.828	10.047
	(0.342)	(0.257)	(0.293)	(0.970)	(0.480)	(0.426)
Geologic Quality (X_i)	0.051	0.077	0.150	0.044	0.031	0.027
	(0.084)	(0.056)	(0.061)	(0.011)	(0.007)	(0.006)
Geologic Quality Squared	0.960	0.340	-0.320	-0.014	-0.010	-0.010
($\times 10^2$)	(0.549)	(0.388)	(0.362)	(0.005)	(0.003)	(0.002)
Residualized Lease Start	-0.439			-0.599		
(months, $\times 10^2$)	(0.044)			(0.162)		
Control Variable ($\hat{V}_i$)		-1.371	-0.466		-1.855	-0.319
		(0.658)	(0.207)		(1.058)	(0.381)
Control Variable Squared		0.848	0.268		0.878	0.083
		(0.441)	(0.226)		(0.915)	(0.417)
Well Order			-0.023			-0.037
			(0.007)			(0.022)
Well Order Squared			0.022			0.155
($\times 10^2$)			(0.037)			(0.198)
Learning Noise Ratio (κ)			4.942			6.140
			(0.923)			(3.392)
Time Splines ($\omega(t)$)	Yes	Yes	Yes	Yes	Yes	Yes
Number of Observations	10,172	7,150	26,854	1,915	1,294	3,573

Notes: The residualized lease start is constructed by regressing the lease start date on a spatial grid (20-by-20 for North Dakota and 12-by-12 for Louisiana) and extracting the residuals. The time (natural cubic) spline basis is constructed using 10 knots placed at equally spaced percentiles (deciles) of the empirical cumulative distribution function of drilling dates across all wells, with boundary knots set at the minimum and maximum observed drilling dates. Standard errors (in parentheses) are calculated from 200 township-level cluster bootstrap replications.

Figure 8 demonstrates that when we correct for this positive selection, the output per well estimate of the time trend $\omega(t)$ for the first wells exhibits steeper growth over time relative to a naive estimate that ignores the declining unobserved quality of marginal drilling sites. As shown in table 4, accounting for selection substantially increases the estimated average annual growth in first-well output from 0.09 to 0.15 log points in North Dakota, and from 0.13 to 0.21 log points in Louisiana.

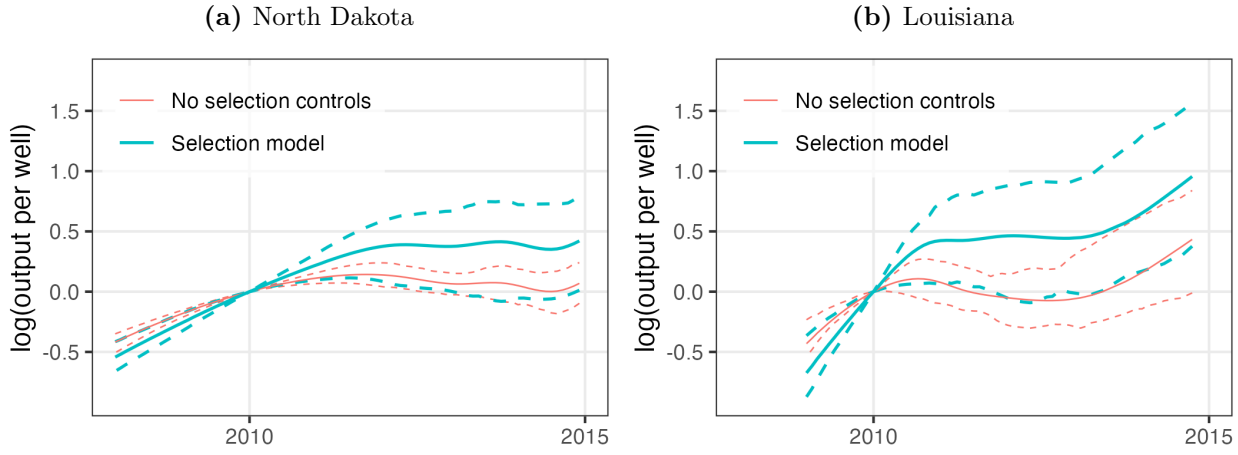
Table 4 illustrates how the growth estimates behave across several alternative specifications. Estimating the model without the observed geologic quality variable (row "Without

Figure 7: Shape of the Control Function



Notes: The range of the plot is restricted to the 2.5th through the 97.5th percentile of the predicted drilling probability data. The control function values are normalized to zero at a predicted drilling probability of 0.5. Dashed lines indicate 95% confidence intervals calculated from 200 township-level cluster bootstrap replications.

Figure 8: Output per Well Trend Estimates from the First Wells



Notes: The plotted range for North Dakota is from the beginning of 2008 to the end of 2015, and for Louisiana from the beginning of 2009 to the end of 2015, reflecting that most first wells are drilled within these periods. The log output per well values are normalized to zero at the beginning of 2010. Dashed lines indicate 95% confidence intervals calculated from 200 township-level cluster bootstrap replications.

quality variable") yields similar growth estimates, which is reassuring because if the assumptions of our model hold, the model should be able to account for selection on the quality measures X_i when those measures are not included in the specification. Dropping the X_i does appreciably inflate the standard errors, however, underscoring the value of including these observed quality measures to improve precision. The estimated growth also remains

Table 4: Average Annual Growth in Log Output per Well

	North Dakota		Louisiana	
Specification	First Well	All Wells	First Well	All Wells
Selection (+ learning) model	0.152	0.121	0.213	0.143
	(0.034)	(0.005)	(0.032)	(0.009)
No learning model		0.122		0.140
		(0.004)		(0.009)
No selection controls	0.087	0.114	0.134	0.131
	(0.008)	(0.004)	(0.011)	(0.008)
Without quality variable	0.153	0.124	0.242	0.144
	(0.051)	(0.005)	(0.092)	(0.009)
Flexible control function	0.145	0.122	0.228	0.141
	(0.044)	(0.005)	(0.032)	(0.009)

Notes: Growth rates are calculated by taking the sum product of the estimated time splines for each well and regressing the resulting aggregate production trend on time (measured in years). No selection controls indicate the model without selection on the first well drilling and learning from previously drilled wells for All Wells. Standard errors (in parentheses) are derived from 200 township-level cluster bootstrap replications.

robust under a more flexible specification that expands the control function to a fourth-order polynomial (row "Flexible control function").

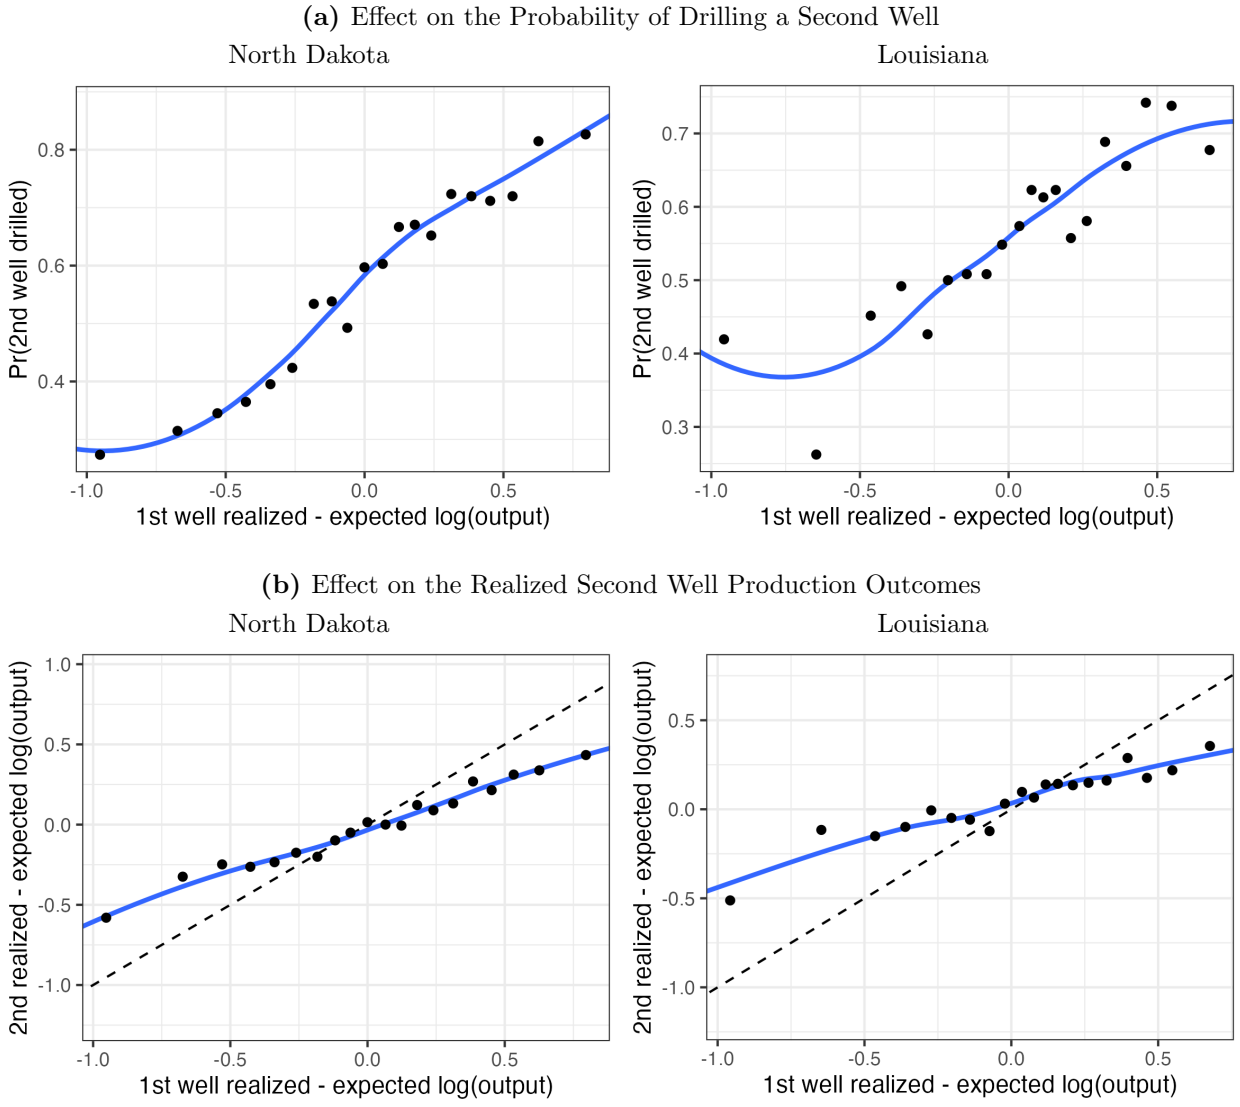
4.2 Estimates using all wells, accounting for learning

Results that use data just from the first well in each section are limiting because few such wells exist in our data after 2015. We therefore employ information from wells subsequently drilled in the same section, while accounting for the possibility that firms learn about the section's geology from observing the output of previously drilled wells using the full model described in section 3.2.

We begin by presenting evidence that realized production outcomes affect firms' subsequent drilling decisions. Panel (a) of figure 9 shows that the probability of drilling a second well in a section increases with the first well's production shock, defined as realized log production minus expected production from the estimated model (including the control function term). Panel (b) shows that the second well's output, relative to expected second-well output from the estimated model (including the control function term but without the learning

term), also increases with the first well's production shock, consistent with the notion that the first well shock contains a signal about the section's true geologic quality. At the same time, the fit in panel (b) falls well short of the 45-degree line, indicating that well-to-well output is mean-reverting, consistent with the idea that each well's output is noisy, and thus that firms cannot perfectly learn the section's true quality from the first well's output.

Figure 9: Effect of the Realized First Well Production Outcomes



Notes: The filled circles display a binned scatter plot, constructed by dividing the sample into vigintiles (20 equal-sized bins) based on the x-axis variable. The solid lines represent LOESS curves fitted to the underlying, unbinned observation-level data. The dashed lines in Panel (b) indicate the 45-degree line.

Given this evidence of learning, we account for it by estimating equation (8), which also allows for effects of depletion and well interference as multiple wells are drilled in the

same section. The full model then accounts for five mechanisms that might affect variation in aggregate output per well over time: (1) observed geologic quality X_i; (2) ex-ante unobserved geologic quality η_i captured by the control function $H(\hat{V}_i)$; (3) learning from previous wells' output captured by the κ-weighted average of previous wells' output shocks; (4) depletion and interference effects captured by the well-order fixed effects λ_j; and (5) the remaining variation $\omega(t)$ that captures changes in technology and operational decisions.

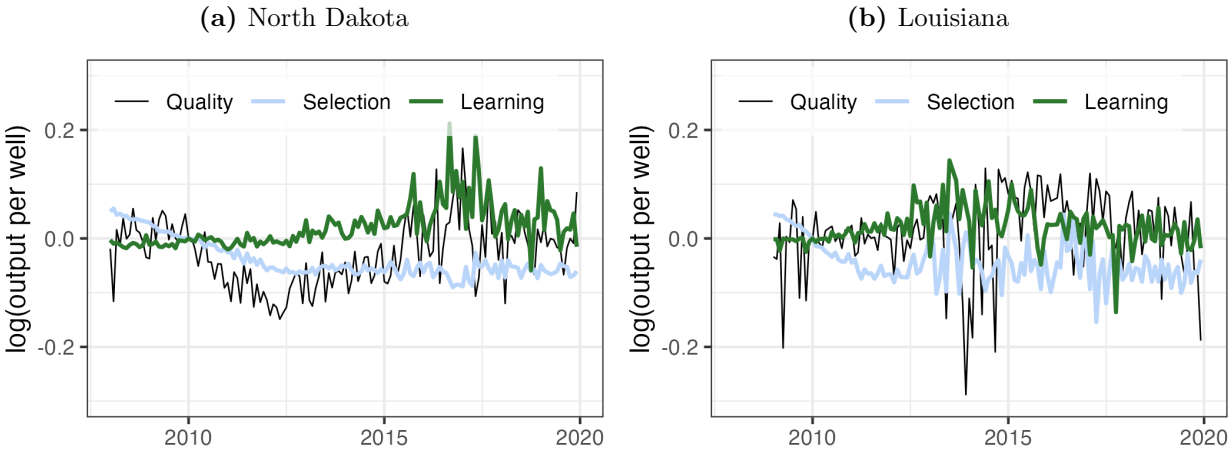
Column "All Wells" in table 3 tabulates the parameter estimates. Consistent with depletion and well interference, the estimated well-order effect is linearly negative, indicating that the output of subsequent wells in a section steadily declines as more wells are added. The estimated learning signal-to-noise ratio κ of 5 for North Dakota and 6 for Louisiana implies that the average output shock from 5 or 6 previously drilled wells carries a weight of 0.5 in the output equation, confirming that operators meaningfully update their expectations based on past drilling outcomes. The visual inspection of the shape of the control function in figure 7 reveals that the positive selection on unobserved quality is weaker than that estimated from the first wells alone. However, the positive selection is estimated with much tighter confidence intervals, reinforcing the conclusion that firms systematically target more promising sites early in the development process.

Our estimates allow us to decompose variation in output per well over time into these mechanisms. Figure 10 plots the contributions of selection on observed quality (OOIP in the Bakken, OGIP in the Haynesille), selection on unobserved quality, and selection on learning from previous outcomes, which together constitute the effects of site selection.[14] In North Dakota, the observable site quality declined initially before rising steadily from 2012 through 2016. Unobserved quality declined steadily from the early sample to 2012, by 0.11 log points, and this low unobserved quality persisted after 2012, in contrast to the observed quality. Over time, learning from the outcomes of previously drilled wells eventually leads firms to select better sites, contributing an additional 0.04 log points to well production. In Louisiana, unobserved quality followed a similar trend to North Dakota, sitting 0.10 log points higher early in the sample than after 2012, while the effect of selection on observed quality does not trend meaningfully over time. Learning from past drilling outcomes led Louisiana firms to select better sites post-2012, driving a 0.04 log-point increase in production.

Figure 11 then plots the sum of the contributions of the three site-selection mechanisms, along with depletion effects λ_j and residual variation $\omega(t)$. In North Dakota, the aggregate effect of the site-selection mechanisms shows a decline of 0.20 log points through 2012,

[14] For each calendar time t, we calculate the contribution of each mechanism using our parameter estimates and the set of wells drilled at t. For instance, we compute the contribution of selection on observed quality as $\bar{X}_i{}'\beta$, where $\bar{X}_i$ is the average observed quality for wells drilled at t.

Figure 10: Output per Well Decomposition of the Site Selection Mechanisms

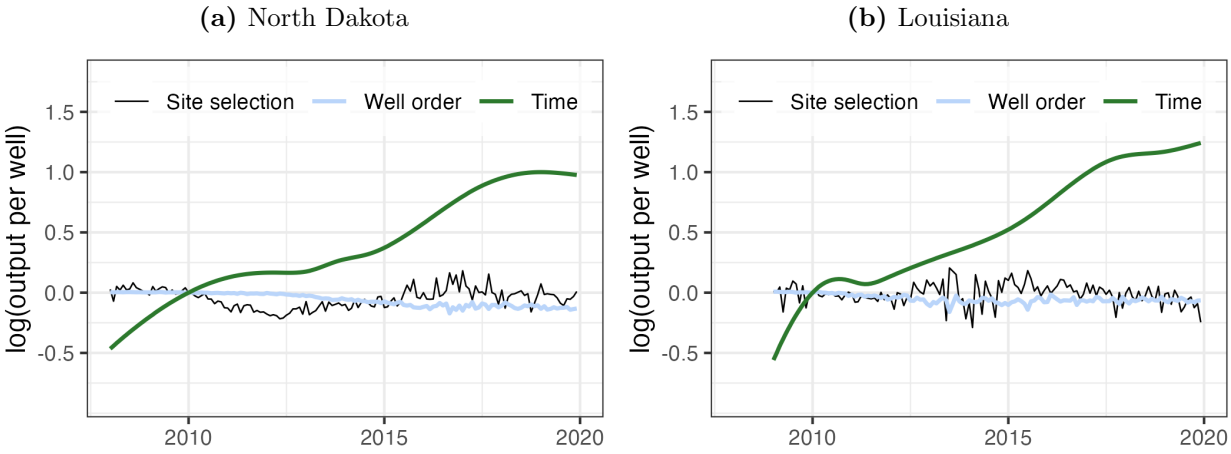


Notes: The plotted range for North Dakota is from the beginning of 2008, and for Louisiana, from the beginning of 2009, reflecting that most wells are drilled after these periods. The log output per well values are normalized to zero at the beginning of 2010.

followed by a bounce back, as improvements in observed quality and learning offset the persistent drag of worse unobserved quality. Additionally, well interference and depletion effects began to occur in North Dakota after 2012, amounting to a penalty of around 0.12 log points. In Louisiana, aggregate site-selection mechanisms have little net effect on production, as unobserved quality and learning largely offset one another, while well interference and depletion effects mirror those in North Dakota (a decline amounting to 0.05 log points). Notably, this figure shows that even though site selection and depletion effects are non-trivial, they are dwarfed by the growth in output per well captured by $\omega(t)$, which we attribute to changes in technology and operational decisions. In North Dakota, $\omega(t)$ increases by 0.6 log points from the early period to 2011, stays relatively flat until 2014, and rapidly grows an additional 0.8 log points through 2018. In Louisiana, $\omega(t)$ follows a similarly dramatic trajectory, rapidly increasing by 0.7 log points from the early period to 2010, and growing an additional 1.1 log points from 2012 to 2019.

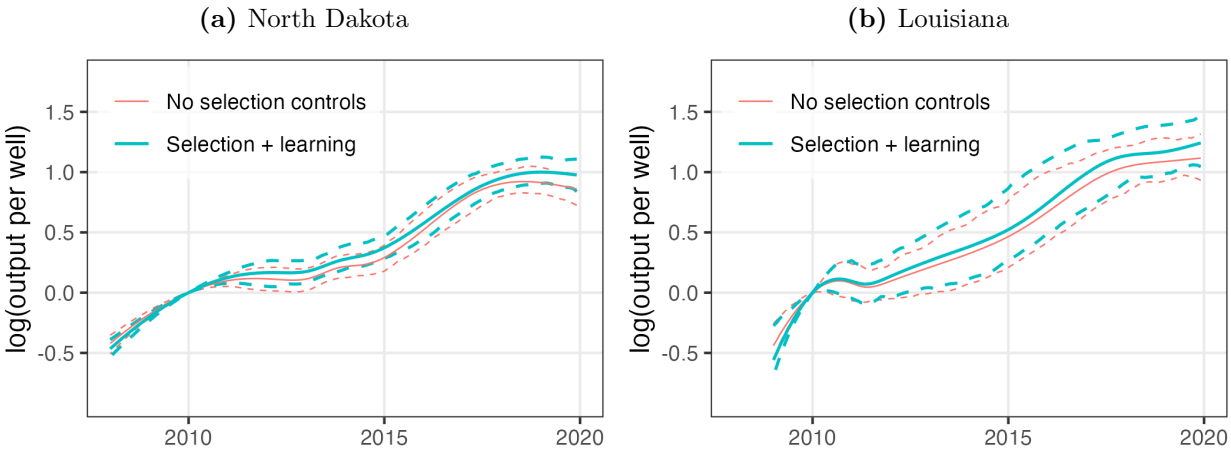
Finally, figure 12 compares our estimated $\omega(t)$ from the full model to an estimate of changes in output per well over time that accounts only for selection on the observable X_i and depletion effects. Overall, accounting for unobserved site selection increases the estimated growth in output over time. This result reflects the fact that positive selection on firms' ex-ante signals of geologic quality out-weighs the effects of learning. These results are also reflected in table 4, which shows that the estimated average annual growth in output per well, using information from all wells in our sample, is about 7% and 9% higher for North Dakota and Louisiana, respectively, when selection is accounted for (comparing the

Figure 11: Overall Output per Well Decomposition



Notes: The plotted range for North Dakota is from the beginning of 2008, and for Louisiana, from the beginning of 2009, reflecting that most wells are drilled after these periods. The log output per well values are normalized to zero at the beginning of 2010.

Figure 12: Output per Well Trend Estimates from the Full Model



Notes: No selection controls indicate the model without selection on the first well drilling and learning from previously drilled wells. The plotted range for North Dakota is from the beginning of 2008, and for Louisiana, from the beginning of 2009, reflecting that most wells are drilled after these periods. The log output per well values are normalized to zero at the beginning of 2010. Dashed lines indicate 95% confidence intervals calculated from 200 township-level cluster bootstrap replications.

row "Selection (+ learning) model" to "No selection controls").

According to table 4, the estimated growth rate does not change meaningfully when the model is estimated without accounting for learning from previously drilled wells (row "No learning model"). In this specification, the well fixed effects λ_j capture both depletion effects and learning, even if they only capture learning imperfectly. The growth estimates

remain highly robust both when excluding the observed geologic quality variable (row "Without quality variable") and when employing a more flexible, fourth-order polynomial control function (row "Flexible control function").

5 What explains the large increase in output per well?

We have found that while shale oil and gas drillers engaged in meaningful site selection, the effects of that selection do not account for the large increase in output per well observed over time, and actually point in the opposite direction. In this section, we unpack the increase in production that is unaccounted for by site selection.

5.1 Selection of productive firms vs. within-firm improvements

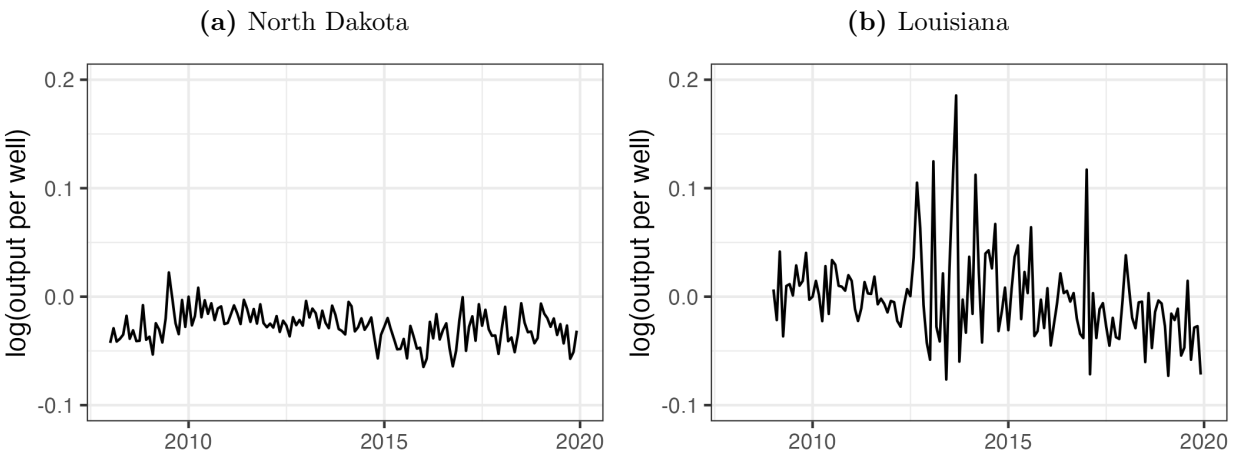
We investigate whether changes in the allocation of drilling activity across firms over time may have played a role in the observed increase in output per well. To understand the potential role of firm-level selection, we estimate our drilling hazard model using first wells and our output model using all wells, including firm fixed effects (FEs) in both to control for time-invariant, firm-specific drilling and operational advantages.[15,16] In our specification, we include firm dummies for the top four firms in each state (Table 2), with the remaining fringe firms pooled into a reference group.

Figure 13 plots the contribution of firm FEs in the output per well decomposition. In North Dakota, the estimated firm FEs exhibit a slight negative trend, dropping by 0.03 log points from 2010 to 2015. In Louisiana, the firm FEs also decline by 0.03 log points, on average, pre and post 2017. The effects of the site selection mechanisms remain largely unchanged in both states relative to the model without firm FEs. Moreover, the output trend estimates for both states do not meaningfully depend on whether firm FEs are included or not, implying that the large increase in well-level output is not driven by an improved selection of more productive firms over time. Instead, the gains are primarily attributable to operational choices within firms. This within-firm production growth contrasts with consolidation or reallocation-driven aggregate productivity improvements, as has been documented in other industries (e.g. Olley and Pakes (1996) in telecommunications equipment and Collard-Wexler and De Loecker (2015) in steel-making).

[15]To estimate the hazard model with firm FEs, we need to assign firms to undrilled sections. To do so, we estimate a classifier based on each drilled section's location (latitude and longitude), drilling date, and shares of leases allocated to various lessees. We then use the classifier to predict the operator for each undrilled section in the sample.

[16]No firm in our data, including Chesapeake in Louisiana, drilled enough wells during our sample to allow us to precisely estimate firm-specific trends in output per well.

Figure 13: Output per Well Decomposition of the Firm Fixed Effects



Notes: The plotted range for North Dakota is from the beginning of 2008, and for Louisiana, from the beginning of 2009, reflecting that most wells are drilled after these periods. The log output per well values are normalized to zero at the beginning of 2010.

5.2 Changes in input use

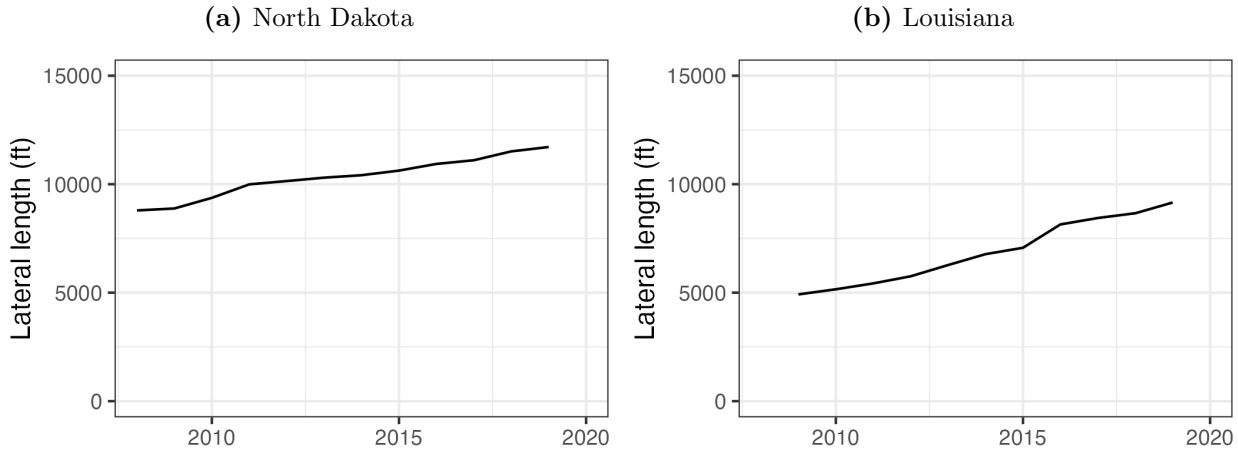
Our decomposition so far divides output into site quality and operational choices, where operational choices encompass both changes in inputs and changes in how effectively those inputs are used. The productivity literature typically distinguishes these two mechanisms by estimating total factor productivity (TFP)—the unobservable proportional increase in output per unit of inputs—often under the assumption that factor productivity is separately identifiable (Syverson, 2011; De Loecker and Syverson, 2021).

Implementing a similar exercise in this setting is challenging because we observe relatively little about input use at the well level. The main exception is well lateral length, which has risen substantially over time (figure 14). The average lateral length of wells has steadily grown from 9,000 feet to 12,000 feet from 2010 to 2019 in North Dakota, and from 5,000 feet to 9,000 feet in Louisiana over the same period.

Including lateral length as an additional covariate in the output model attenuates the estimated time trend, reflecting that overall output growth is partially driven by the industry-wide shift toward longer laterals (figure 15). The average annual growth in log output per well drops by 13% in North Dakota, and by 16% in Louisiana with lateral length controls. However, even controlling for well length, we continue to find a large increase in output per well, indicating that the growth in lateral length alone does not account for the productivity gains.

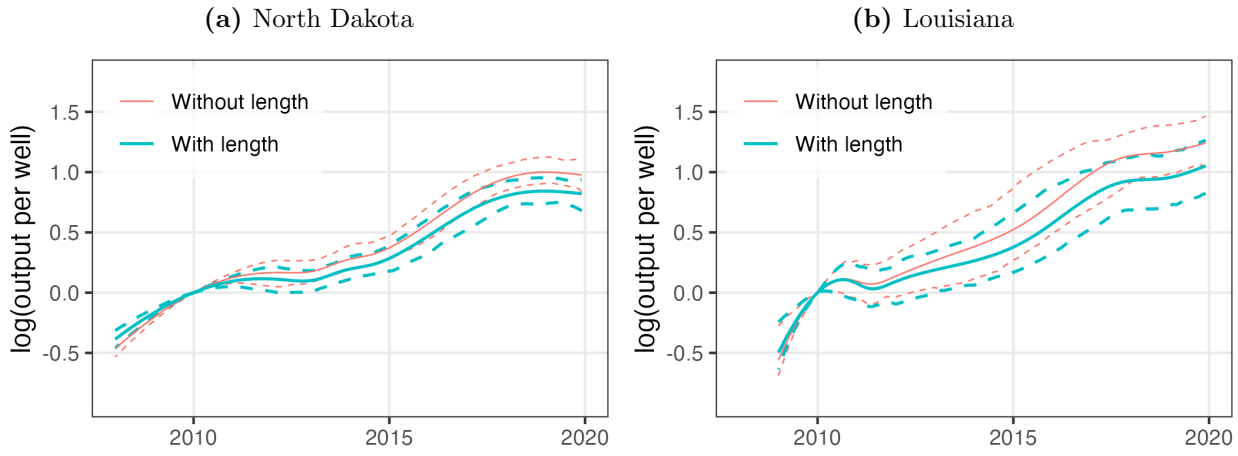
For a subset of wells, we also observe proppant and water use from Novi Labs (2023), and we find that these inputs have grown even faster than length: firms are applying more

Figure 14: Trends in Well Lateral Length



Notes: Figure shows the average annual well lateral length calculated from our analysis sample. The plotted range for North Dakota is from 2008, and for Louisiana, from 2009, reflecting that most wells are drilled after these periods.

Figure 15: Output per Well Trend Estimates with Lateral Length Controls

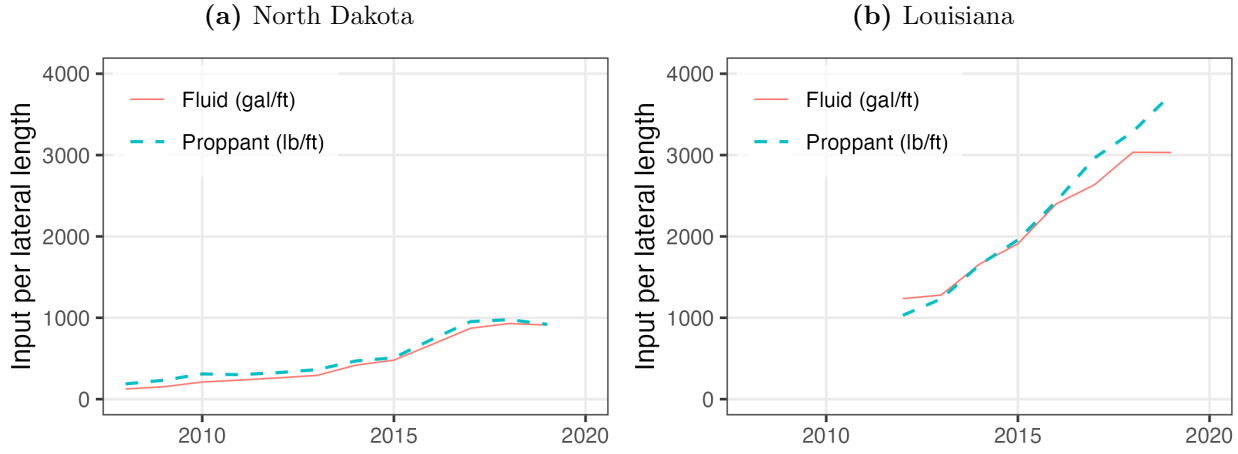


Notes: The plotted range for North Dakota is from the beginning of 2008, and for Louisiana, from the beginning of 2009, reflecting that most wells are drilled after these periods. The log output per well values are normalized to zero at the beginning of 2010. Dashed lines indicate 95% confidence intervals calculated from 200 township-level cluster bootstrap replications.

inputs per foot drilled (figure 16). Between 2013 and 2018, average fluid and proppant use per foot in North Dakota grew from 300 gal/ft and 400 lb/ft to 900 gal/ft and 1,000 lb/ft. In Louisiana, this growth was even more pronounced, with average usage per foot increasing from 1,300 gal/ft and 1,200 lb/ft to 3,000 gal/ft and 3,300 lb/ft over the same period. These large increases in input use are clearly correlated over time with the large increases in output per well that we estimate after accounting for site selection.

A natural question is then whether our estimated increase in output per well is driven entirely by increased input use, holding TFP fixed, or by improvements in firms' ability or application of technology to combine these inputs together. While it is tempting to estimate a traditional production function — in which each well's output were written as a function of observed inputs inputs plus an unobserved TFP term — using the subset of wells for which all inputs are observed, we think doing so would be an error. One problem is econometric: we lack a credible set of instruments that would shift each input, leaving other inputs fixed. A more fundamental conceptual problem is that there exists a multitude of additional costly inputs — such as chemical additives, pumping equipment, downhole equipment, and pumping pressures applied — that we do not observe and are likely correlated with the observed inputs.[17] Thus, even if we were able to credibly identify a traditional production function, labeling the disturbance term as TFP would not be justified. To make progress, we therefore turn in the next sub-section to data on prices and the total cost of drilling and completing wells.

Figure 16: Proppant and Water Use Trends



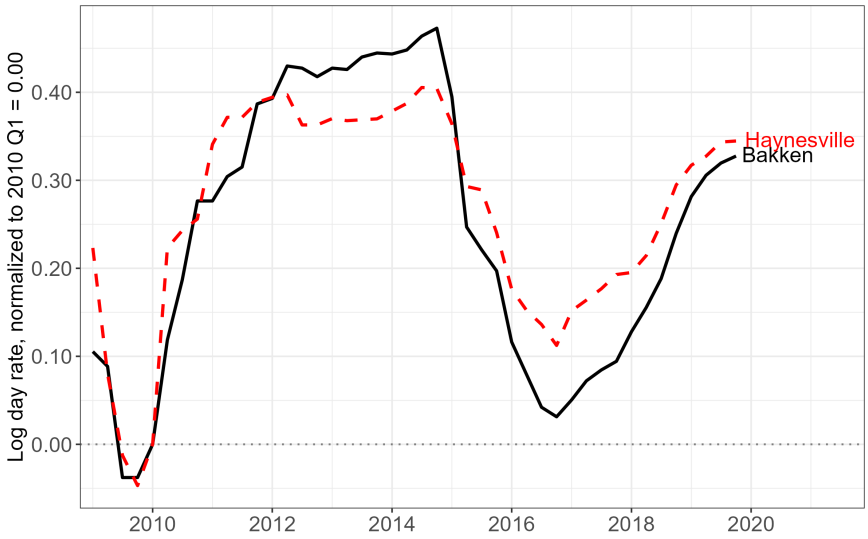
Notes: Figure shows the average annual input use calculated from Novilabs data. The plotted range for North Dakota is from 2008 to 2019. For Louisiana, the range is from 2012 to 2019, since fewer than 10 wells were recorded each year before 2012.

5.3 Changes in costs and input prices

Why did input use increase so much? One possibility is that firms learned over time that large volumes of inputs improved output and profits, even holding unobserved technological factors

[17]Moreover, there exist a large variety of different frac sands that operators use, but we only observe a single measure of total weight of proppant applied.

Figure 17: Rig rental day rates



Notes: Figure shows quarterly day rate data from Enverus. We use Enverus's Rocky Mountain region for the Bakken, and "ArkLATex" for the Haynesville. Data correspond to depth ratings of 10,000–12,999 feet before 2018, and horsepower class 1000–1499hp from 2018 onward. Log day rates are normalized to zero in the first quarter of 2010.

constant, per the mechanism described in Covert (2015). Another is that firms adopted new productivity-increasing technolgies, such as equipment enabling multi-stage fracturing, that are complementary to input use but not directly observable to the econometrician. Either of these mechanisms could appropriately be described as an increase in productivity. But other explanations exist as well. Input use may have increased simply because the cost of inputs decreased, which would represent not a productivity increase but rather just a rational response to a cost decrease. Finally, another possible (though unsatisfying) explanation could be that firms were making a mistake by increasing input use: while they increased output they did so at the expense of profits, with no real increase in productivity.

To investigate these mechanisms, we have compiled time series data on prices for key inputs, as well as data on drilling and completion costs for a subset of wells in our sample. The most important cost item for an E&P company drilling and completing a shale well is the rental of drilling and completion rigs, along with associated crews and equipment. These rental costs, which exceed $10,000 per day throughout our sample, scale with the total length of the well and the intensity of the frac job. To assess how costs per day have changed over time, we gathered data from Enverus on average rental rates for drilling rigs. Time series of rental rates for both the Bakken and the Haynesville are plotted in figure 17. Day rates rise and fall with drilling activity, consistent with the idea from Anderson, Kellogg and Salant (2018) that rigs and crews are scarce in the short-to-medium run, so that their supply curve

is upward-sloping. Rates are also strongly correlated across plays, consistent with the fact that rigs are mobile.

Day rates fall by roughly 0.4 log points during 2015–2016, following the large drop in oil prices at the end of 2014 (recall figure 3) and the consequent decrease in drilling activity. But as drilling activity recovers starting in 2017, especially in the Bakken, day rates recover and by end-2019 are only about 0.1 log point lower than they were at their 2014 peak.[18]

We also investigated relevant producer price indices (PPIs) that are available in the St. Louis Federal Reserve Bank's FRED database. The PPI for hydraulic fracturing sand decreased by about 30% during 2015–2016 before modestly rebounding. And the PPI for "oil country tubular goods" decreased by about 20% during 2015 before rebounding by about 10% during 2016 (after which the PPI is no longer available).

These input price decreases on their own would increase input use even in the absence of any productivity improvement. However, they did not occur in isolation: oil and gas prices more than halved during 2014–2016. The fact that the fall in output prices exceeded the fall in input prices implies that firms' observed increase in input use was not driven by price changes alone, and if anything occurred despite price changes that should have decreased rather than increased input use.[19]

We next study data on wells' drilling and completion costs. For the Bakken, cost data are available for a subset of wells that were involved in regulatory proceedings regarding payments to mineral owners that did not participate in leasing.[20] In the Haynesville, cost data are available for most wells because firms must report drilling and completion costs to the Louisiana Department of Natural Resources in order to qualify for severance tax relief (Herrnstadt et al., 2024). These data are plotted in figure 18. In North Dakota, drilling and completion costs rise during the 2010–2014 drilling boom, and then fall by about 10% during 2014–2015 before leveling off. In Louisiana, costs fall by 10–20% from 2012–2015, and then in 2018–2019 return to their 2011 level.

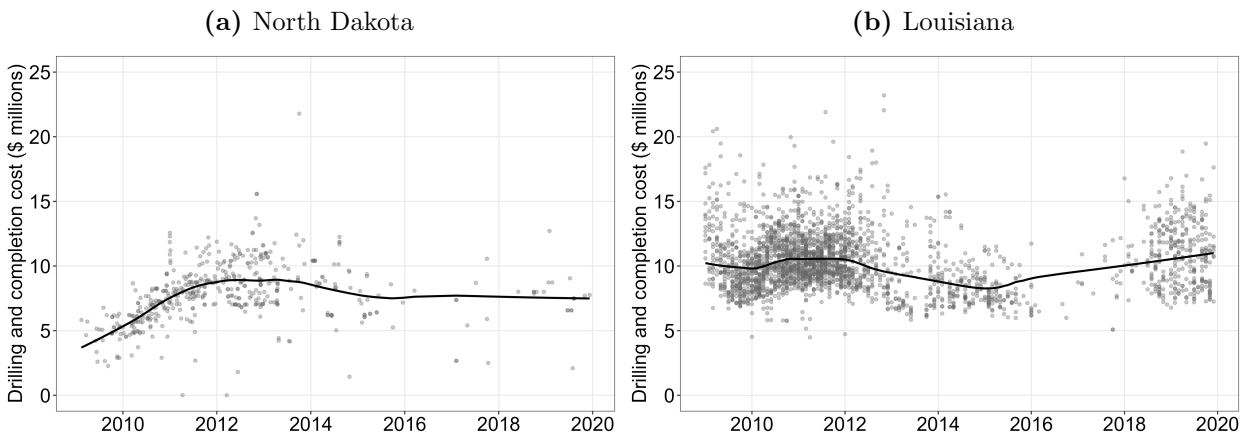
Overall, our analysis of inputs, prices, and costs has found that during our sample period: (1) well lateral length increased by 30–80%, and sand and water inputs per length more than doubled; (2) oil and gas prices more than halved during 2014–2016; (3) input prices fell by

[18]Shale oil drilling also increased substantially in the Permian Basin (West Texas and New Mexico) during 2015–2020.

[19]For example, if oil production is Cobb-Douglas in inputs, then an input's use will be an increasing function of that ratio of the output price to the input's price.

[20]When one or more mineral owners in a section does not sign a lease and does not "participate" in a well as a working interest owner, North Dakota allows the operating firm to impose a "risk penalty" on these owners, which reduces their share of revenue from the well's production. The risk penalty is based on the cost of drilling and completing the well, and as part of the regulatory process the firm must report its internal cost estimate for the well by providing its internal "authorization for expenditure" document. These cost data were also used in Covert (2015).

Figure 18: Reported drilling and completion costs per well



Notes: Figures show reported drilling and completion costs for a subsample of wells for which data are available. Each point is an individual well, and the solid line is a LOWESS curve fitted to the underlying data. Cost data are available and plotted for 509 wells in North Dakota and 2697 wells in Louisiana.

20–40% during 2015–2016, with some rebound afterwards; and (4) reported drilling and completion costs track input price changes, though with a dampened magnitude. The price data imply that firms' increased use of inputs was not simply a response to price changes. And the reported cost data imply that the large increase in input use, and the consequent increase in hydrocarbon output, was not associated with a large increase in the total cost of drilling and completing shale wells. The absence of a meaningful cost increase during our sample period implies that firms were not just using more and more inputs: they were using them more productively. The data do not allow us to discern whether this productivity increase arose because firms learned about the production function over time, a la Covert (2015), or because of the adoption of new technologies that are complementary to input use (and not directly observed by us). But either mechanism is consistent with the same substantive conclusion: the large increase in output per well during the shale boom was primarily driven by improvements in the productivity with which E&P companies brought real resources—sand, water, rigs, crews, and land—together.

6 Conclusion

In this paper, we decompose the increase in output per well during the U.S. shale boom into effects stemming from firms' selection of high versus low-quality sites to drill, and those stemming from changes in firms' operational decisions and factor productivity. We do so by introducing a joint model of output and investment decisions that accounts for endogeneity

of investment timing that stems from firms' private information about potential drilling sites' geologic quality, and from firms' learning about quality from the outcomes of previous investments.

We find evidence for both forms of site selection: among sites that have not yet been drilled, firms tend to first drill higher-quality sites, and then once many sites have been drilled at least once, they tend to preferentially drill sites that had better production outcomes. However, these mechanisms fall far short of explaining, and if anything work modestly against, the large observed growth in output per well during the shale boom. Instead, it appears that over time firms have made meaningful improvements in the productivity with which they bring together inputs to produce output. This large productivity boom suggests that future growth in shale production may be more resilient to the depletion of high-quality sites than if the growth had been primarily driven by site selection. It also suggests that the energy transition may be slower than what one might conclude by simply looking at the decline in clean energy costs alone.

A fruitful line of future inquiry would be to investigate the timing of the shale productivity boom. Curiously, we find that the rate of increase in output per well, especially in the Bakken, was largest during 2015–2017, which was a period of low oil prices and low drilling activity. This timing is consistent with the "slack time" theory of innovation from the management literature (Nohria and Gulati, 1996). A standard model of a firm would predict that firms should invest more in output-enhancing technology when prices are high than when prices are low, since the return to additional production is greater when the output price is high. The "slack time" theory, however, argues that an opposite mechanism may be important: during bust periods, oil and gas companies scale back their drilling activity but retain technical employees on payroll in order to have them available again when prices rise. They will thus have excess human capital, leaving those employees with "slack time" to analyze geologic and engineering issues they had postponed during the boom. This research and development activity then leads to higher productivity, both during the bust and persisting when the next boom arrives. Our findings are consistent with this theory, and with trade press that argued that downturns could be an opportunity for the industry to become more productive.[21] That said, our results are not dispositive, and future work could investigate the theory more directly should data become available on firms' technical employment and R&D activities over time.

Finally, while our paper focuses on the evolution of productivity in the shale oil and gas

[21]During the low-price environment of 2015–2016, almost one quarter of all issues of the Society of Petroleum Engineers' *Journal of Petroleum Technology* contained editorials or guest articles arguing that the downturn was an opportunity to 'use data more effectively" and 'learn" to make better engineering decisions. For examples, see Rassenfoss and Henni (2015); Betz (2015); Jacobs (2016a,b); Rassenfoss (2016).

industry, our approach for decomposing output growth into site selection versus operational decisions is applicable to a range of other settings. There is a natural parallel to other forms of energy resource development where output is site-dependent, such as geothermal, wind, and solar projects (though for wind and solar, there is arguably little private information about resource quality). Beyond energy, our approach could be applied to agricultural settings, where farmers select which plots of land to plant and then make operational decisions about how to manage those plots; franchising, where the success of a franchisee depends on both the quality of the location where it opens and how efficiently each location is run; or private equity, where the success of an investment depends on both the quality of the target firm and the operational decisions made by the private equity owner.

References

Acemoglu, Daron, Philippe Aghion, Lint Barrage, and David Hémous, "Climate Change, Directed Innovation, and Energy Transition: The Long-run Consequences of the Shale Gas Revolution," September 2023. NBER working paper 31657.

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

Amin Gherabati, S., John Browning, Frank Male, Scott Hamlin, Katie Smye, Mark Walsh, Svetlana A. Ikonnikova, Guinevere McDaid, and Casee Lemons, "Evaluating hydrocarbon-in-place and recovery factor in a hybrid petroleum system: Case of Bakken and three forks in North Dakota," in "SPE/AAPG/SEG Unconventional Resources Technology Conference 2017" 2017.

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

Baumeister, Christiane and Lutz Kilian, "Understanding the Decline in the Price of Oil since June 2014," *Journal of the Association of Environmental and Resource Economists*, 2016, *3*, 131–158.

Betz, Jack, "Low Oil Prices Increase Value of Big Data in Fracturing," *Journal of Petroleum Technology*, 2015, *67* (04), 60–61.

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).

Casero, Alberto, "20 Years of Horizontal Multistage Completions: A Summary of Industry Evolution in Unconventional and Conventional Plays," *Abu Dhabi International Petroleum Exhibition & Conference Proceedings*, 2021, p. D031S069R004.

Collard-Wexler, Allan and Jan De Loecker, "Reallocation and Technology: Evidence from the US Steel Industry," *American Economic Review*, January 2015, *105* (1), 131–71.

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 versus Informal Negotiations in Texas Oil and Gas Leasing," *American Economic Review*, March 2023, *113* (3), 628–63.

Curtis, Trisha, "Unravelling the US Shale Productivity Gains," 2016. Oxford Institute for Energy Studies, OIES Paper: WPM 69.

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).

Herfindahl, Orris C., "Depletion and Economic Theory," in Mason Gaffney, ed., *Extractive Resources and Taxation*, University of Wisconsin Press, 1967.

Herrnstadt, Evan, Ryan Kellogg, and Eric Lewis, "Drilling Deadlines and Oil and Gas Development," *Econometrica*, 2024, *92* (1), 29–60.

Hodgson, Charles, "Information Externalities, Free Riding, and Optimal Exploration in the UK Oil Industry," Working Paper 33067, National Bureau of Economic Research October 2024.

Hotelling, Harold, "The Economics of Exhaustible Resources," *Journal of Political Economy*, 1931, *39* (2), 137–175.

Imbens, Guido W and Whitney K Newey, "Identification and estimation of triangular simultaneous equations models without additivity," *Econometrica*, 2009, *77* (5), 1481–1512.

Jacobs, Trent, "Downturn represents stress test for unconventional hydraulic fracture modeling," *Journal of Petroleum Technology*, 2016, *68* (02), 44–49.

_ , "Shale Completions Improve, Despite Shrinking Number of New Wells," *Journal of Petroleum Technology*, 2016, *68* (04), 30–32.

Lazear, Edward P., "The Peter Principle: A Theory of Decline," *Journal of Political Economy*, 2004, *112* (S1), S141–S163.

Leonard, Bryan and Dominic P. Parker, "Fragmented Ownership and Natural Resource Use: Evidence from the Bakken," *Economic Journal*, 2020, *131*, 1215–1249.

Loecker, Jan De and Chad Syverson, "An industrial organization perspective on productivity," in "Handbook of industrial organization," Vol. 4, Elsevier, 2021, pp. 141–223.

Newey, Whitney K., "Two-Step Series Estimation of Sample Selection Models," *The Econometrics Journal*, 2009, *12* (S1), S217–S229.

Nohria, Nitin and Ranjay Gulati, "Is slack good or bad for innovation?," *Academy of management Journal*, 1996, *39* (5), 1245–1264.

Novi Labs, "Upstream Energy Data," accessed 30 May, 2023.

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," 2024. working paper.

Rassenfoss, Stephen, "Pressure Test for E&P Innovation," *Journal of Petroleum Technology*, 2016, *68* (05), 39–43.

_ **and Abdelghani Henni**, "Low oil prices make innovation a priority," *Journal of Petroleum Technology*, 2015, *67* (02), 56–61.

Steck, Andrew, "Industry Dynamics with Social Learning: Evidence from Hydraulic Fracturing," 2022. working paper.

Syverson, Chad, "What determines productivity?," *Journal of Economic literature*, 2011, *49* (2), 326–365.

U.S. Energy Information Administration, "United States Produces More Oil Than Any Country, Ever," 2024. [Online; `https://www.eia.gov/todayinenergy/detail.php?id=61545`; accessed 11 March, 2024].

Wethe, David, "U.S. Oil Companies Use AI For Faster Shale Drilling, Mitigating OPEC Production Cuts," *World Oil*, March 2024.