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

Thomas R. Covert Konan Hara

Ryan Kellogg Richard L. Sweeney*

July 7, 2026

Abstract

The goal of this paper is to understand why output per well increased so dramatically during the shale boom. 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. But after accounting for selection, as well as reservoir depletion, firms' operational improvements emerge as the largest driver of output growth. These improvements occurred within firms, rather than being driven by reallocation of drilling activity to more productive firms. Price and cost data suggest that these gains were driven by improvements in the productivity with which firms brought together drilling inputs.

*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

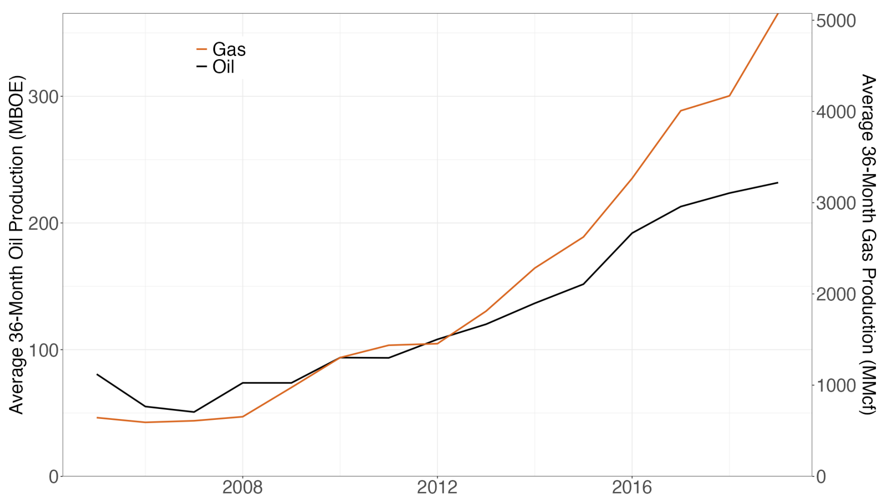

Notes: Output from the first 36 months of production for wells drilled in each year, across all U.S. shale plays. See section 2 for details on data sources and construction.

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. Measuring drilling productivity conditional on site quality is challenging because 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). The direction of site selection is thus theoretically ambiguous, and potentially changes over time. This ambiguity has implications for the future trajectory of shale production. If firms initially knew a great deal about geology, such that Hotelling forces dominate, the history of growth in shale wells' output could continue for some time into the future, since firms have demonstrated an ability to draw more and more output out of less productive geologic locations. Conversely, if the learning story dominates, such that firms increased output per well largely by converging on the best locations, then the future of shale production may be more limited, since these locations are inherently scarce.

To estimate output 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 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 given 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 account for 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.

We find evidence of positive selection on unobserved quality for first wells, consistent with firms using private geological information to prioritize more promising drilling opportunities early on. Accounting for this selection increases our estimates of output growth, conditional on site quality, among first wells during the early years of the shale boom in both states, 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 previously 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 of learning in both states that we study: firms are more likely to drill a second well in a section if the first well's output is higher than expected. In addition, we find that higher than expected first-well output also predicts higher second-well output, but the magnitude of this effect is mean-reverting. This mean reversion implies that observed output from the first well is part signal and part noise, so that learning is incomplete.[1]

To account for the dynamic selection introduced by firm learning, we estimate a Bayesian updating model in which expected output and the firm's drilling decision are influenced by

[1] As we discuss in section 3.3, incomplete learning invalidates a simpler empirical strategy of simply using a model of well-level output that controlled for spatial variation in geology using granular section fixed effects.

previously drilled wells' output surprises, relative to expectation. This model amounts to including a second control function in the outcome equation that accounts for the possibility that the output surprise for the jth well in a given section will be correlated with the surprises from wells 1 through $j-1$.[2] Unlike Agerton (2020), we account for the fact that well-to-well production outcomes are highly noisy, and allow learning to be incomplete. The identifying assumption in this model is that after the first well is drilled, the only factors that jointly affect subsequent wells' drilling decisions are the outcomes of a correctly specified set of previously drilled wells. We first consider a model where the firm only learns from wells drilled within the same section, and then expand the learning set to include wells drilled in other nearby sections. We estimate this learning model jointly with the first-well outcome equation (including the leasing-based control function), using GMM.

We find that firms learn meaningfully from realized production outcomes. In the model in which firms only learn from wells drilled within the same section, we find that improved site selection increased wells' output by 0.12 log points in North Dakota from 2012 to 2017, and by 0.06 log points in Louisiana from 2013 to 2015. When we expand the learning set to include wells drilled in nearby sections, we find that the effect of learning is larger, increasing to 0.41 log points in North Dakota and 0.14 log points in Louisiana. Still, these effects from learning are considerably smaller than that found for Louisiana in Agerton (2020), which assumed that firms learn completely about geologic quality after drilling each section's first well and then optimally choose locations for subsequent wells based on this information.[3]

Our model also accounts for the fact that wells' output must eventually decline as more and more wells are drilled within a single section, due to reservoir depletion and well interference effects. We address this mechanism in our output equation by accounting for the number of wells drilled in each section. We find that by the end of our sample period (end-2019), the average well's output suffers a depletion and well interference penalty of -0.20 log points in North Dakota and -0.12 log points in Louisiana.

While the site selection and depletion forces that we estimate have economically meaningful magnitudes, they nonetheless jointly account for less than half of the overall increase in output per well over time. Conditional on site quality and depletion, we find that the average growth in output was 0.083 log points per year in North Dakota and 0.079 log points per year in Louisiana. Nearly all of this growth occurred in two periods: the early days of the

[2] These surprises are defined relative to the firm's expectation of output before the first well was drilled. We define the surprises precisely in section 3.3.

[3] Unlike our decomposition 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. Agerton (2020) also only models learning from wells drilled within the same section, not learning from wells drilled in nearby sections.

boom before 2010, and the three-year period immediately following the late-2014 collapse in oil prices.

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 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 decreases 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.

The evidence thus 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 wells' oil and gas production, conditional on the quality of the sites drilled, was driven primarily by productivity improvements, and that these improvements were complementary to increases in well-level input use. Overall, both these productivity improvements and changes in site selection contributed to the observed increase in output per well during the shale boom, with the productivity improvements being the larger factor.

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, depletion effects, and 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 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-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

[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 2: Map of Major Shale Plays

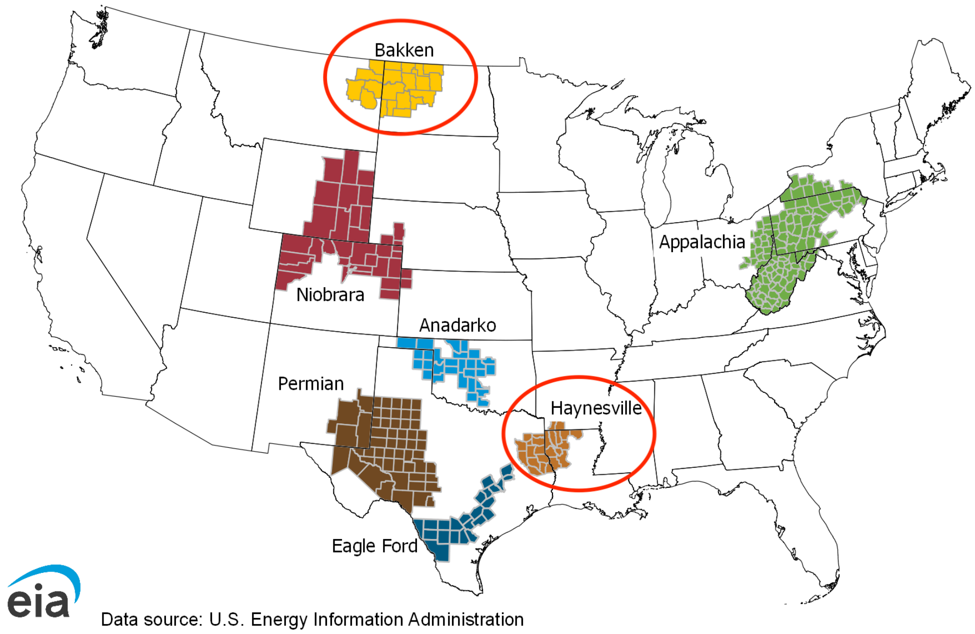

Source: U.S. Energy Information Administration.

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.

One partial driver of increased output per well is simply that wells have gotten longer over time (figure 4). 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.

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

Figure 3: Price, Completions, and Production Over Time

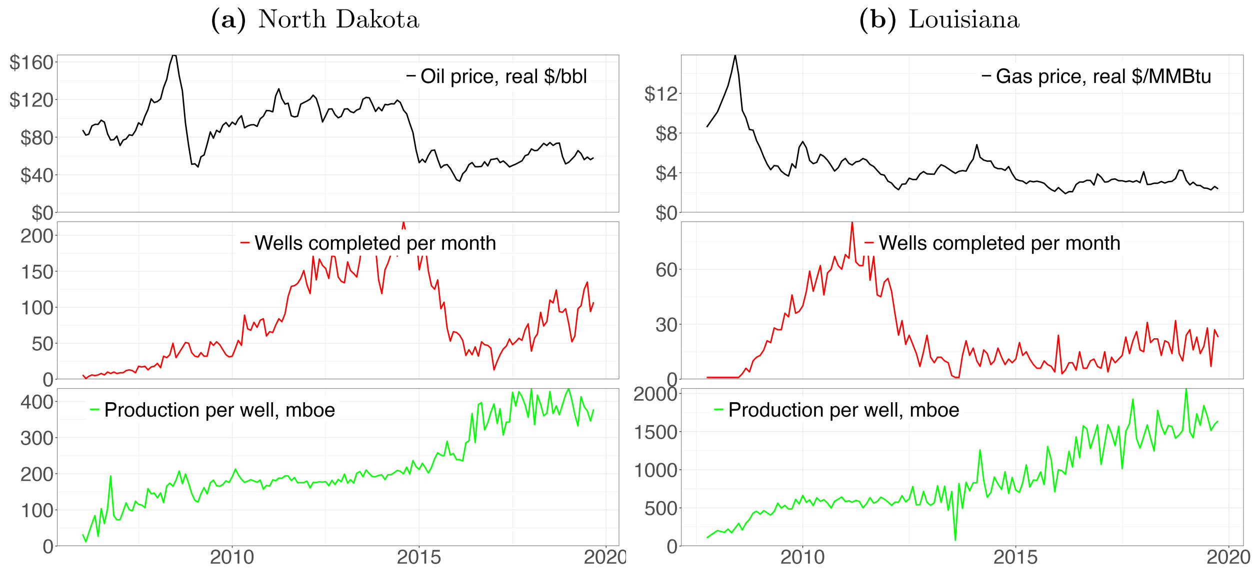

Notes: The top panel shows average real oil prices (West Texas Intermediate) and average real natural gas prices (Henry Hub) over time. The middle panel shows completions per month in North Dakota and Louisiana. The bottom panel shows total monthly production for new wells by month. See section 2 for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

Figure 4: Trends in Well Lateral Length

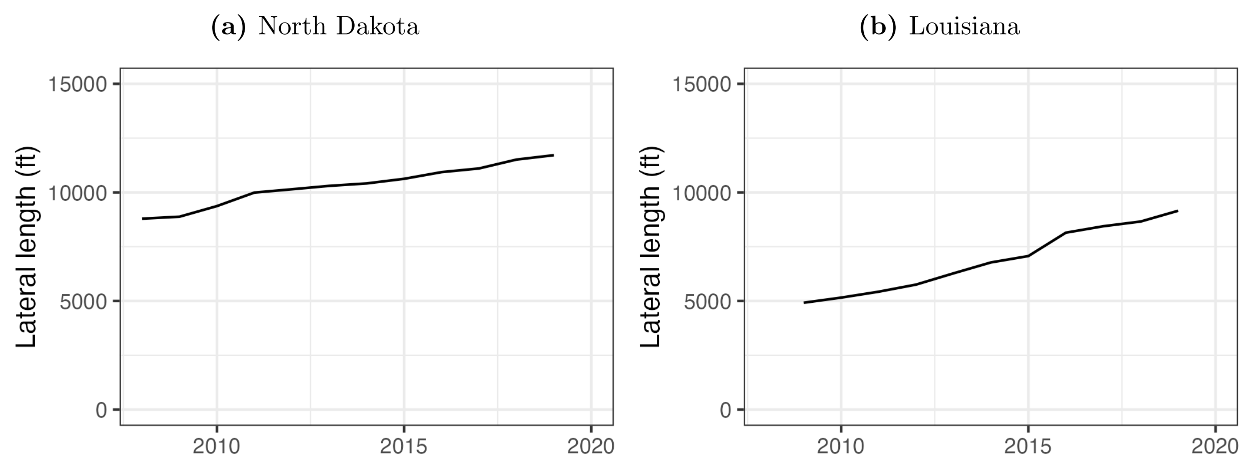

Notes: Figure shows the average annual well lateral length calculated from our analysis sample. The plotted range for North Dakota is from 2008, and for Louisiana, from 2009, reflecting that most wells are drilled after these periods.

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, 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 5 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 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 deduplication, 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

Figure 5: Example Section and Parcel Map

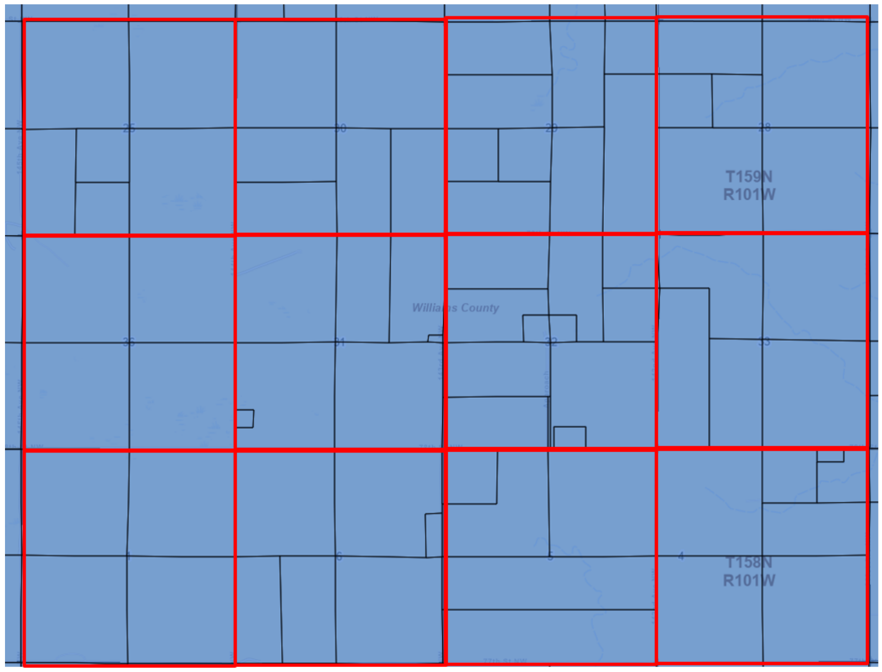

Notes: Parcel map from the North Dakota GISHUB. The red boxes indicate sections, and the black lines indicate the boundaries of individual parcels that are owned by different mineral owners.

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 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.

[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.

Figure 6: Distribution of when sections become actively leased and are first drilled

(a) North Dakota **(b)** Louisiana

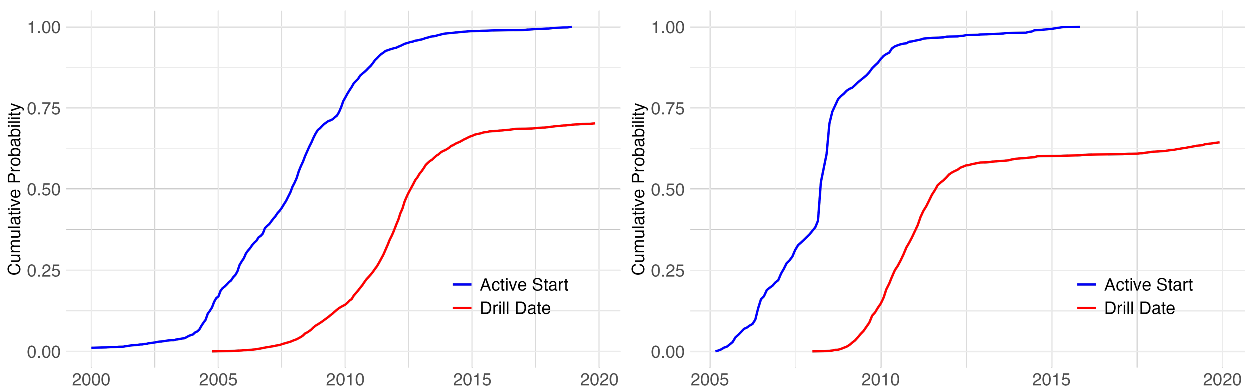

Notes: Figure shows the cumulative distribution of the date a section if 50% leased and the date of the first well drilled.

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 6 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 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.

Many sections in our sample are drilled more than once. Table 1 presents our sample's summary statistics and shows that conditional on any drilling, the average section has 3.76 wells in the Bakken and 2.75 wells in the Haynesville. In both plays, the distribution of wells per section is substantially right-skewed. Figure 7 plots annual time series, for each state, of the number of "first wells" drilled in each section and the number of "later wells" drilled in sections that have already been drilled. First well drilling peaks in 2012 in the Bakken and 2010 in the Haynesville, and in both states the number of later wells drilled exceeds the number of first wells drilled after 2012. Few first wells are drilled in either state after 2014.

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

[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.

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.

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 8 plots the within-year

[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.

Figure 7: First Wells vs Later Wells

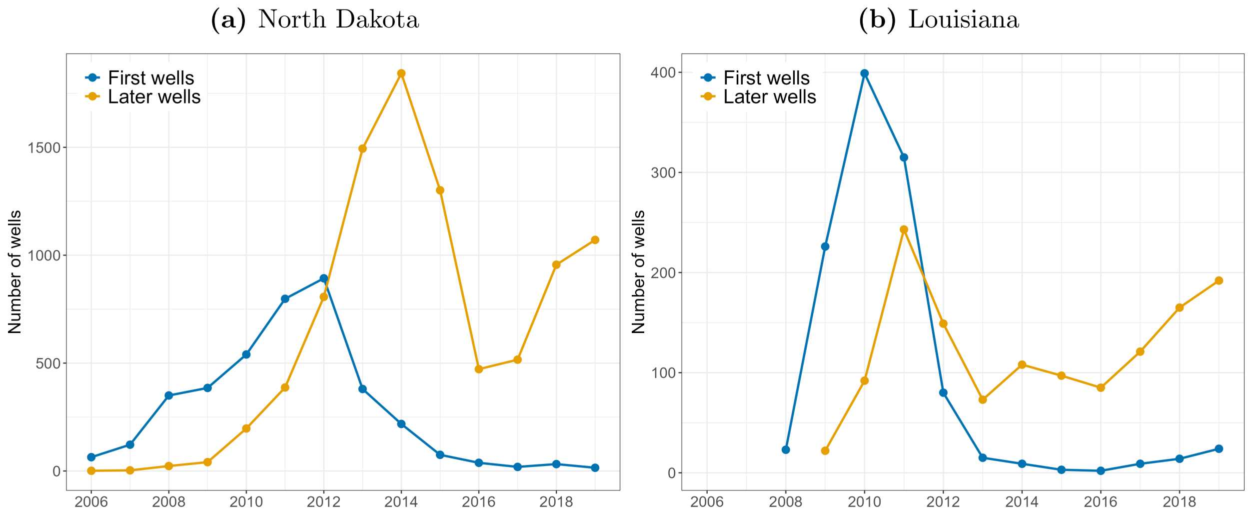

Notes: For each state, the figure shows the number of first and later wells drilled each year. First wells are the first well drilled in a section, and later wells are any subsequent wells drilled in that section. The first well time series corresponds to the drill timing cdf from figure 6.

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.

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%

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.

Figure 8: Herfindahl-Hirschman Index (HHI) of Drilling Firms over Time

(a) North Dakota **(b)** Louisiana

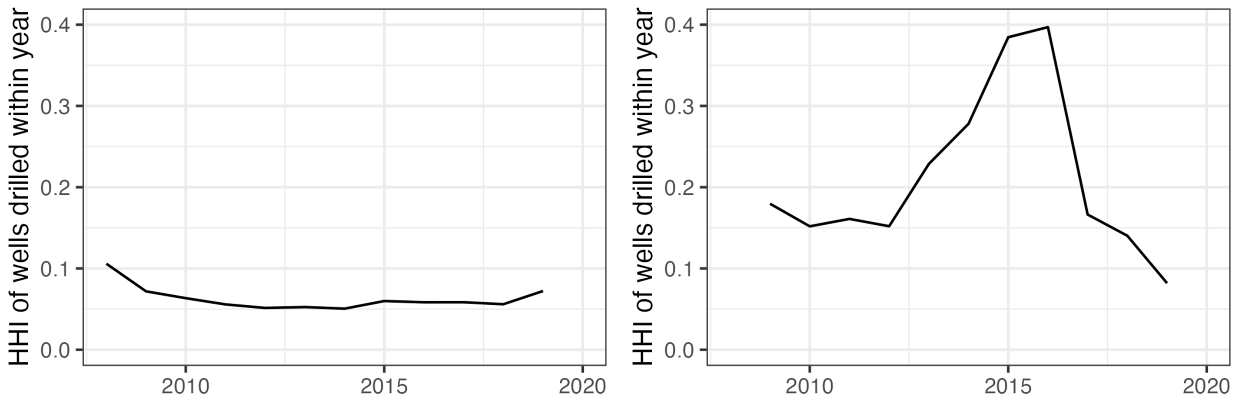

Notes: The plotted range for North Dakota is from 2008, and for Louisiana, from 2009, reflecting that most wells are drilled after these years.

3 Model and estimation

This section presents the econometric model that we use to identify and estimate the contributions of site selection to changes in output per well over time. Sub-section 3.1 begins by presenting our modeling framework and outcome equation. Sub-section 3.2 then presents our model of selection into whether and when a drilling opportunity (henceforth, a "section") receives its first well. Sub-section 3.3 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. Finally, sub-section 3.4 discusses how we estimate the model.

3.1 Outcome equation

Our outcome variable of interest is each well's log barrels of oil equivalent (boe) per unit length of its horizontal lateral. We study output per unit length, rather than simply output, because we are interested in how drilling investment opportunities—i.e. potential reservoir targets—are selected and executed. Dividing output by lateral length accounts for the fact that a well with a two-mile lateral exhausts twice as much reservoir area as a well with a one-mile lateral.

Let y_{ij} denote the log boe per unit length produced by the j-th well drilled in section i.

We model y_{ij} as:

$$y_{ij} \;=\; \omega(t_j(i)) + X_i'\beta + \underbrace{\eta_i + \theta_i}_{\substack{\text{unobserved}\\\text{site quality}}} + \lambda(j) + \varepsilon_{ij}, \tag{1}$$

where X_i, η_i, and θ_i represent observed and unobserved measures of site quality. $\omega(t_j(i))$ represents the extent to which firms' choices increase output conditional on quality, and we model its evolution over time as a flexible function of the date $t_j(i)$ that the well was drilled. $\lambda(j)$ is a function of well order j that accounts for reservoir depletion and well interference as more and more wells are drilled into the same section (we normalize $\lambda(1) = 0$). Finally, ε_{ij} is a mean-zero idiosyncratic output shock that is realized after the well is completed.

We are interested in separately identifying the evolution of $\omega(t)$, which reflects operational choices and technological advances that increase output per well, and the evolution of site quality. The identification challenge comes from unobserved quality, which we break into two terms that facilitate how the model captures two different forms of site selection. η_i denotes the firm's unbiased belief about the section's quality prior to drilling the first well, and by construction it is mean zero conditional on X_i. θ_i is the difference between section i's true geologic quality and the firm's ex ante belief η_i, and we will allow firms to learn about θ_i by observing previously drilled wells' production outcomes. By construction, θ_i is mean zero conditional on X_i and η_i. Sub-sections 3.2 and 3.3 below discuss our strategies for capturing selection on η_i and firms' beliefs about θ_i, respectively.

Equation (1) does not include terms that account for inputs into drilling and completing wells, such as the length of the horizontal lateral or the volume of water and sand used during hydraulic fracturing. We omit inputs because we observe relatively few of them—lateral length for all wells, and sand and water use for a subset—and because these inputs are likely to be correlated with the use of other costly inputs that we do not observe, such as equipment use, the pumping pressure applied, chemical additives, and so on. Nor do we have a credible instrument for shifting input use, either individually or even as a bundle of observed and unobserved inputs.[8] The time and site effects in equation (1) should therefore be interpreted as a composite of the direct effects of time and location on output with indirect effects stemming from the correlation of input use with time and location. Appendix A.1 formalizes the intuition for this interpretation via a model in which inputs are a mediator for the effects of time and location on output. Later in section 5 we will use data on input use, prices, and costs to shed light on the drivers of the time effects $\omega(t)$ that we estimate in equation (1).

[8] We also cannot apply methods such as those discussed in Ackerberg, Caves and Frazer (2015) because our data do not have the panel structure typical of papers in the production function literature that study annual data for manufacturing plants.

3.2 First-well selection and timing

In this sub-section, we focus on the first well $j = 1$ drilled in each section. 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 $\omega(t)$. (The estimate of β is also likely to be attenuated, per the logic in Olley and Pakes (1996) and Herrnstadt et al. (2024).)

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 captures 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 $\omega(t)$ (up to an additive constant) by writing the expected log output of the first well in each section i, conditional on $t_1(i)$, X_i, and Z_i as:

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

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 $t_1(i) = g\big(X_i, Z_i, u_i\big)$:

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

Our model has a parallel to binary selection models in the spirit of Heckman (1979), but here we model not just *whether* but also *when* each section i is drilled. If we were just modeling selection into drilling (regardless of timing), then instead of specifying and estimating a hazard (and thus survival) model, we would instead estimate a model for the probability of drilling, i.e. some function $f(X_i, Z_i) \to (0,1)$, and define $V_i = 1 - f(X_i, Z_i)$ for all drilled sections.[9] The need to use a selection equation that explicitly accounts for time—as in our hazard / survival model—is driven by our objective of accounting for the impact of time on output via the $\omega(t)$ term in the output equation (1).

Identification of our joint output and survival model follows Imbens and Newey (2009) and requires 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.

Assumption 2 and the independence of u_i are ensured to hold by the structure of the survival model in expression (4). The independence of η_i requires that Z_i be uncorrelated with unobserved section quality, an assumption we discuss in detail in section 4.1. Independence is a stronger assumption than orthogonality and is not testable, though we note that this assumption is made implicity in prior work that models oil and gas firms' drilling decisions and well-level output (Bhattacharya et al., 2018; Agerton, 2020; Herrnstadt et al., 2024; Ordin, 2024). Assumption 3 requires that there be sufficient variation in the instrument Z_i to shift drilling timing across the entire calendar time window of interest for any given value of X_i. This assumption is analagous to (though stronger than) a rank condition from a linear simultaneous equations model.

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.

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).[10] We eschew that approach here because: (1) our primary concern is

[9] We also depart from the canonical Heckman (1979) model by not assuming a parametric joint distribution for (η_i, u_i), and instead model $E[\eta_i|V_i]$ flexibly (see Newey (2009) for a flexible selection correction in standard binary selection models).

[10] 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

measuring how time and location affect output, 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 the 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 estimates to be inconsistent.

3.3 Subsequent-well selection, learning, and depletion effects

As shown in table 1, many sections in our sample are ultimately drilled more than once. These late wells complicate the estimation of time-varying productivity, for two reasons. 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. We account for this effect by including $\lambda(j)$ in the output equation (1). Second, firms may learn something about latent quality θ from the first well's output. Indeed, as we show in section 4.2 below, the probability that a section receives a second well is strongly related to the "surprise" component of the first well's output, consistent with firm learning. This learning is important because the control function developed in the previous section only accounts for selection due to firms initial beliefs η. If firms learn over time which sections are truly high-quality, and then drill those sections relatively frequently, our estimates will overstate the extent to which increases in output per well over time are attributable to the mechanisms underpinning $\omega(t)$.

To see how an estimate of ω in equation (1) 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 is drilled. Assume there are no nearby sections that firms can learn from, so that the firm's beliefs about θ_i are only updated based on the output of the first well in section i. Define e_{i1} as the first well's output surprise: $e_{i1} = y_{i1} - \omega(t_1(i)) - X_i'\beta - E[\eta_i|V_i]$. This output surprise is the unexpected component of output from the firm's perspective, combining the ex post output shock ε_{i1} with the part of θ_i unknown to the firm. We then have the expectation:

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

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 output surprise 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

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.

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 (1) to overstate the increase in productivity over time.[11]

Extending the model to the jth well and learning from wells drilled in nearby sections, let $\mathcal{E}_{ij}$ denote the history of output surprises (relative to the firm's ex ante expectation before drilling the first well in each section but accounting for the depletion effects $\lambda(j)$) for all previous wells in the section and nearby sections (defined formally in Appendix A.2). The expected output of the jth well, conditional on the history of outputs of all previous wells in the section and nearby sections (collected in $\mathcal{Y}_{ij}$), is

$$E[y_{ij}|t_j(i), X_i, Z_i, \mathcal{Y}_{ij}] = \omega(t_j(i)) + X_i'\beta + \lambda(j) + E[\eta_i|V_i] + E[\theta_i|\mathcal{E}_{ij}]. \tag{5}$$

To consistently estimate equation (5), we address learning by conditioning on previous wells' output surprises, thereby accounting for the $E[\theta_i|\mathcal{E}_{ij}]$ term. The identifying assumption is that after drilling begins in and around section i, firms only update their beliefs about section i's latent productivity based on the output from a (correctly specified) set of previously drilled wells. We consider two sets of prior wells to learn from. First, we assume that firms only learn about θ_i from wells drilled in section i ("within section" learning). Second, we broaden this set to include wells drilled in other sections close to section i ("across section" learning).[12]

Finally, we note that the existence of incomplete learning invalidates an alternative empirical strategy that is tempting by virtue of its simplicity. Suppose that we attempted to estimate output growth conditional on site selection 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 infor-

[11] If the $E[\theta_i|e_{i1}]$ term is not accounted for, the well order function $\lambda(j)$ will capture some of the learning effect, causing $\lambda(j)$ to increase with j relative to the case when this function captures 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 surprise from the first well, not just whether or not a subsequent well was drilled.

[12] Agerton's (2020) also studies productivity and learning in the Haynesville shale. Our approach differs for two reasons. First, that paper imposes 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. Second, that paper only allows firms to learn from previous wells within the same section, whereas we allow production outcomes from nearby sections to be informative about section i's quality.

mation 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 output growth conditional on site.[13] This mechanism for bias in a simple section fixed effect regression an analogue to the mechanism described in Lazear (2004) to explain declines in employee performance following promotions (also known as the "Peter principle").

3.4 Estimation

We estimate the hazard $h(t, X_i, Z_i)$ 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. We define the instrument Z_i based on information on when section i was leased, as discussed in detail in section 4.1. 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, $H(\hat{V}_i) = H(\hat{F}(t_1(i)|X_i, Z_i))$, as an additional regressor to account 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}. If there are no nearby sections that firms can learn from ("within-section learning model"), then the posterior mean of θ_i before drilling the jth well is

$$E[\theta_i | e_{i1}, \ldots, e_{i,j-1}] = \underbrace{\frac{j-1}{(j-1) + \sigma_s^2/\sigma_p^2}}_{\text{Signal precision relative to prior}} \times \underbrace{\frac{1}{j-1}\sum_{k=1}^{j-1} e_{ik}}_{\text{Average output surprise}} . \tag{6}$$

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 parameter $\kappa = \sigma_s^2/\sigma_p^2$.[14]

[13] When we estimate a model with section fixed effects, we find that the average annual growth in $\omega(t)$ is only 0.058 log points per year in ND and 0.064 log points per year in LA. These are smaller than the estimates we obtain using our full model of 0.096 in ND and 0.081 in LA (table 4).

[14] We can allow for a more general learning process (e.g., non-normal distributions or non-Bayesian learning)

Extending this idea to allow learning from nearby sections, we also assume that the latent quality of nearby section i', $\theta_{i'}$, is correlated with section i's latent quality θ_i, which allows output surprises from wells in nearby sections to be informative about θ_i. The correlation or informativeness of wells in nearby section i' about section i's quality presumably becomes weaker as the distance between section i and nearby section i' increases. Thus, we parameterize the correlation between θ_i and $\theta_{i'}$, denoted as $\rho_{d(i,i')} \in [0,1]$, as a function of the discrete distance $d(i,i') \in \{0,1,\ldots,D\}$ between sections i and i', where $d(i,i') = 0$ if and only if $i' = i$, and $\rho_0 = 1$ by definition.

Let $\bar{e}_{i'}^{ij}$ and $J_{i'}^{ij}$ be the average output surprise and the number of wells in section i', respectively, before drilling the jth well in section i. Then, a multivariate normal Bayesian learning model implies that the posterior mean can be written as the precision-weighted average of the average output surprises from prior wells in each section that firms learn from (derivation in Appendix A.2):

$$E[\theta_i|\mathcal{E}_{ij}] = \sum_{i' \in \{i' : \rho_{d(i,i')} \neq 0\}} \frac{\pi_{i'}^{ij}/\rho_{d(i,i')}}{1 + \sum_\ell \pi_\ell^{ij}} \times \bar{e}_{i'}^{ij}, \text{ where } \pi_{i'}^{ij} = \frac{\rho_{d(i,i')}^2 J_{i'}^{ij}}{(1-\rho_{d(i,i')}^2) J_{i'}^{ij} + \kappa}. \tag{7}$$

The formula represents the key features of Bayesian updating: section-level signals ($\bar{e}_{i'}^{ij}$) contribute proportionally to their precision $\pi_{i'}^{ij}$, repeated wells within a section exhibit diminishing marginal informational value through $J_{i'}^{ij}$, non-perfect correlation ($\rho_{d(i,i')} < 1$) reduces informational content, and the prior enters as a unit precision term (because of the precision normalization by σ_p^2). If $\rho_{d(i,i')} = 0$, and thus $\pi_{i'}^{ij} = 0$, for all nearby sections $i' \neq i$, then firms only learn from wells drilled in the same section, and the posterior mean reduces to equation (6).

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 $\hat{\bar{e}}_{i'}^{ij}$ denote the average residual from all wells in section i' before drilling the jth well in section i. We then have the following estimation equation for the jth well in section i:

$$y_{ij} = \omega(t_j(i)) + X_i'\beta + H(\hat{V}_i) + \lambda(j) + \sum_{i' \in \{i' : \rho_{d(i,i')} \neq 0\}} \frac{\pi_{i'}^{ij}/\rho_{d(i,i')}}{1 + \sum_\ell \pi_\ell^{ij}} \cdot \hat{\bar{e}}_{i'}^{ij} + \varepsilon_{ij}. \tag{8}$$

---

and still keep the estimation tractable by assuming that the average output surprise is a sufficient statistic for θ_i before drilling the jth well. Then, $E[\theta_i|e_{i1},\ldots,e_{i,j-1}] = E[\theta_i|\frac{1}{j-1}\sum_{k=1}^{j-1} e_{ik}] = l_{j-1}(\frac{1}{j-1}\sum_{k=1}^{j-1} e_{ik})$, 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. To allow for spatial learning, the learning function would take as input the vector of average output surprises from prior wells in each section that firms learn from.

Equation (8) is nonlinear in its parameters, both due to the functional form for κ and the ρ_d parameters and because $\hat{\bar{e}}^{ij}_{i'}$ is itself a function of all parameters. We therefore estimate it by GMM, as detailed in Appendix A.2, using estimates from the within-section learning model as starting values.

We specialize our estimation to incorporate the following three empirical features. First, firms drill "batches" of wells contemporaneously within a section. We group wells in a section into a single batch if they are drilled within 60 days of the preceding wells. Second, firms need some amount of time after observing production outcomes to incorporate that learning into the drilling operations of a subsequent batch. For each well j in section i, we include prior wells that are drilled at least 60 days before the start of the batch containing well j as wells that firms can learn from and thus contribute to the learning signals in $\mathcal{E}_{ij}$. Thus, all wells drilled within the same batch share the same posterior mean by construction.

Third, a well may appear in multiple sections according to our sample construction procedure in section 2.3. If the same learning well appears in multiple sections in the learning signal calculation, we only keep the section-well $i'j'$ whose section centroid is closest to section i to avoid double-counting the same output surprise as a learning signal for multiple sections. We aggregate these deduplicated wells at the section level for each learning section i' to construct the section-level signal $\hat{\bar{e}}^{ij}_{i'}$.

We use discrete distance up to $D = 3$ based on section-centroid distance from section i: $d(i, i') = 1$ for sections within $(0, 2]$ miles, $d(i, i') = 2$ for sections within $(2, 4]$ miles, and $d(i, i') = 3$ for sections within $(4, 6]$ miles. Sections farther than 6 miles are excluded from the learning wells or, equivalently, assumed to have zero latent quality correlation with section i, $\rho_d = 0$. We also parameterize the latent quality correlations through a single distance decay parameter $\gamma_\rho \geq 0$, as $\rho_d = \exp(-\gamma_\rho \delta_d)$ for $d = 1, 2, 3$, where δ_d is the average distance of all distance d wells that contribute to the learning signals in $\mathcal{E}_{ij}$. This restriction makes the correlation structure consistent with radial distance decay: along a ray from section i, the product of the correlations over two adjacent segments equals the correlation implied by their combined distance.

Finally, we parameterize the well order effects $\lambda(j)$ as a polynomial function of j. For inference, we compute standard errors by cluster bootstrapping the entire estimation procedure at the township (36 section) level, with 200 replications.

4 Results

4.1 Estimates using only the first well in each section

We begin by estimating the output equation, with log boe per unit length as the outcome, using only the first well in each section, accounting for selection into whether and when a section is drilled. 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,[15] 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.[16] Moreover, parcel ownership fragmentation alone will not capture variation across mineral owners in their willingness to lease.

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, as shown in figure 6. 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-

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

[16] 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.

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 (recall the variation in parcel fragmentation across adjacent sections displayed in figure 5), 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 9. 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 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.

Figure 10 demonstrates the impact of this positive selection on the estimated time path of ω. Estimating the model without controlling for selection yields a flatter growth path as it fails to account for the fact that sections drilled later are lower quality than early sections. Once this change in quality is accounted for, the estimated impact of firms' operational decisions looks larger in both states. We summarize the difference in average growth rates in table 4. This table presents the average implied annual change in $\omega(t)$, calculated by taking the sum product of the estimated time splines $\omega(t)$ for each well and regressing the resulting aggregate production trend on time (measured in years).[17] The results show that accounting for selection substantially increases the estimated average annual growth in first-well output from 0.05 to 0.10 log points in North Dakota, and from 0.07 to 0.14 log points in Louisiana.

[17] Although table 4 presents our estimates of a linear average trend, the models are estimated with flexible splines for the functions of time in both the output equation and hazard function.

Table 3: Estimation Results

	North Dakota				Louisiana			
	Hazard	First Well	Within Section	Across Section	Hazard	First Well	Within Section	Across Section
Intercept	-9.539	0.897	0.700	1.013	-11.689	1.209	1.468	1.587
	(0.342)	(0.270)	(0.296)	(0.300)	(0.970)	(0.451)	(0.383)	(0.396)
Geologic Quality (X_i)	0.051	0.058	0.155	0.051	0.044	0.034	0.029	0.029
	(0.084)	(0.055)	(0.058)	(0.061)	(0.011)	(0.007)	(0.006)	(0.006)
Geologic Quality Squared	0.960	0.459	-0.316	0.360	-0.014	-0.011	-0.010	-0.010
($\times 10^2$)	(0.549)	(0.379)	(0.342)	(0.396)	(0.005)	(0.003)	(0.003)	(0.002)
Residualized Lease Start	-0.439				-0.599			
(months, $\times 10^2$)	(0.044)				(0.162)			
Control Variable ($\hat{V}_i$)		-0.875	-0.352	-0.227		-1.833	0.137	0.104
		(0.702)	(0.198)	(0.259)		(1.039)	(0.260)	(0.222)
Control Variable Squared		0.412	0.091	0.259		1.041	-0.411	-0.279
		(0.447)	(0.216)	(0.251)		(0.917)	(0.298)	(0.257)
Well Order			-0.026	-0.051			-0.050	-0.071
			(0.006)	(0.007)			(0.017)	(0.017)
Well Order Squared			0.038	0.148			0.275	0.410
($\times 10^2$)			(0.030)	(0.035)			(0.142)	(0.143)
Learning Parameters								
Signal Noise Ratio (κ)			3.554	7.730			4.407	6.064
			(0.842)	(1.210)			(1.813)	(4.541)
Spatial Correlation Decay (γ_ρ)				0.321				0.933
				(0.128)				(0.312)
Time Splines ($\omega(t)$)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Number of Observations	10,172	7,150	26,854	26,854	1,915	1,294	3,573	3,573

Notes: Column "Hazard" reports estimates of the first-well drilling-timing model parameters. Column "First Well" reports the production model parameters estimated using only the first well drilled in each section. Columns "Within Section" and "Across Section" report the production model parameters estimated using all wells, allowing firms to learn from previously drilled wells within the same section and across nearby sections, respectively. 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.

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

Figure 9: Shape of the Control Function

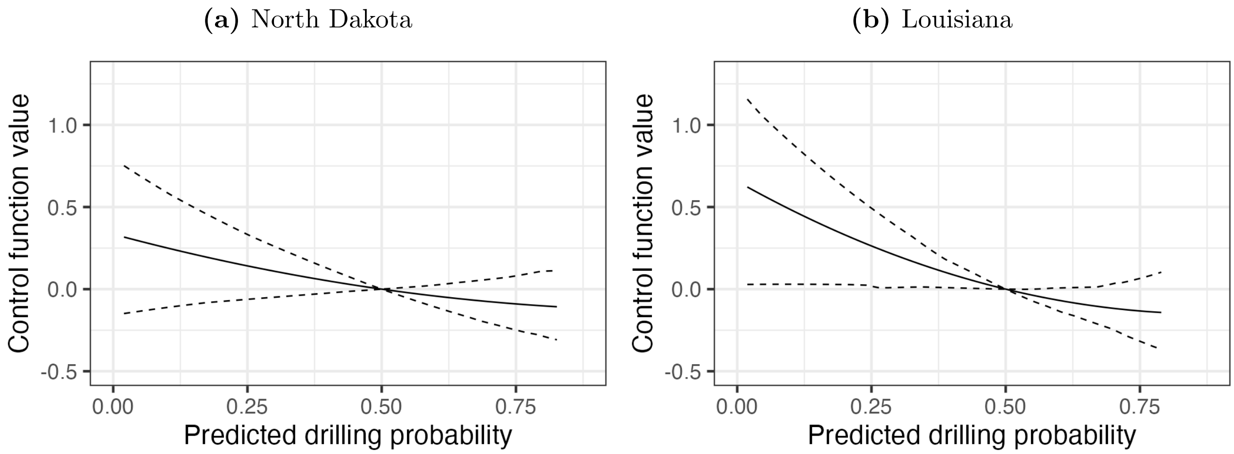

Notes: The range of the plot is restricted to the 2.5th through the 97.5th percentile of the predicted drilling probability data. The control function values are normalized to zero at a predicted drilling probability of 0.5. Dashed lines indicate 95% confidence intervals calculated from 200 township-level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

Figure 10: Output per Lateral Length Trend Estimates from the First Wells

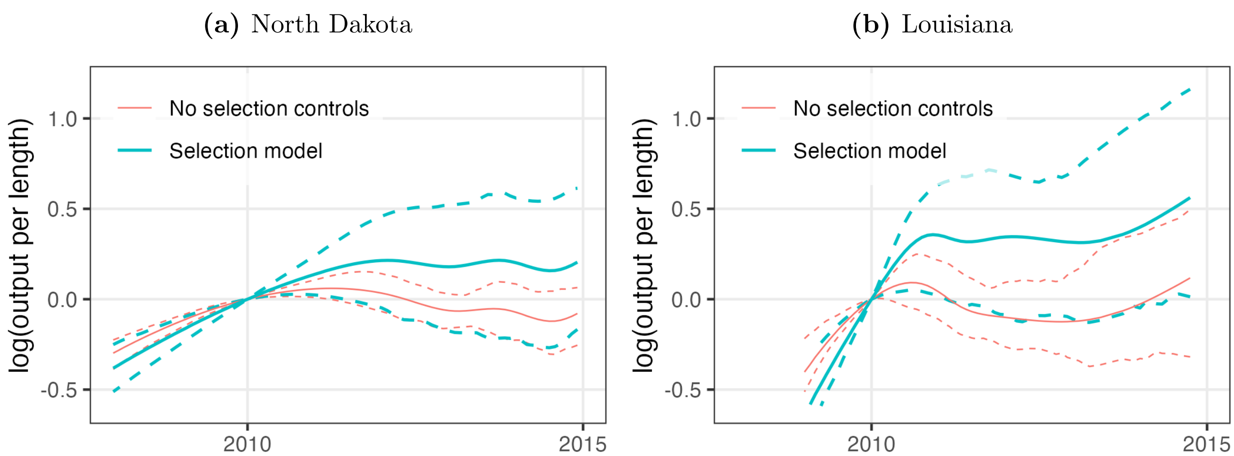

Notes: The plotted range for North Dakota is from the beginning of 2008 to the end of 2015, and for Louisiana from the beginning of 2009 to the end of 2015, reflecting that most first wells are drilled within these periods. The log output per lateral length values are normalized to zero at the beginning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

Table 4: Average Annual Growth in Log Output per Lateral Length

	North Dakota			Louisiana		
Specification	First Well	Within Section	Across Section	First Well	Within Section	Across Section
Selection (+ learning) model	0.096	0.096	0.083	0.143	0.081	0.079
	(0.036)	(0.005)	(0.006)	(0.029)	(0.007)	(0.006)
No learning model		0.095	0.095		0.079	0.079
		(0.004)	(0.004)		(0.007)	(0.007)
No learning, no well order		0.085	0.085		0.072	0.072
		(0.004)	(0.004)		(0.007)	(0.007)
No selection controls		0.086	0.086		0.071	0.071
		(0.005)	(0.005)		(0.007)	(0.007)
No selection, no well order	0.045	0.081	0.081	0.075	0.070	0.070
	(0.009)	(0.004)	(0.004)	(0.010)	(0.007)	(0.007)
Without quality variable	0.101	0.098	0.084	0.208	0.082	0.081
	(0.053)	(0.005)	(0.005)	(0.096)	(0.007)	(0.007)
Flexible control function	0.075	0.096	0.083	0.160	0.080	0.079
	(0.048)	(0.005)	(0.006)	(0.030)	(0.007)	(0.006)

Notes: Columns "First Well", "Within Section", and "Across Section" report average annual growth rates of $\omega(t)$ implied by production models estimated using first wells only, all wells with learning within the same section, and all wells with learning within the same section and across nearby sections, respectively. Growth rates are calculated by taking the sum product of the estimated time splines $\omega(t)$ 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. Well order effects $\lambda(j)$ are included in the "No learning model" and "No seleciton controls" rows. Columns "Within Section" and "Across Section" coincide for specifications that do not involve learning. Standard errors (in parentheses) are derived from 200 township-level cluster bootstrap replications.

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 in that section and nearby sections using the learning model described in section 3.3.

We begin by presenting evidence that realized production outcomes affect firms' subsequent drilling decisions within the same section. Panel (a) of figure 11 shows that the probability of drilling a second well in a section increases with the first well's output surprise, defined as realized log production minus expected production from the estimated within-section learning model (including the control function term). Panel (b) shows that the second well's output, relative to expected second-well output from the estimated within-section learning model (including the control function term but without the learning term), also increases with the first well's output surprise, consistent with the notion that the first well output surprise 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.

Given this evidence of learning, we investigate its contribution to changes over time in wells' output in two steps. First, we estimate the within-section version of equation (8) (using the posterior mean belief for $E[\theta_i|e_{i1}, \ldots, e_{i,j-1}]$ given by equation (6)). Then, we estimate the full model that also allows for learning from wells drilled in nearby sections. Throughout, we also account for the effects of depletion and well interference as multiple wells are drilled in the same section by including a polynomial for $\lambda(j)$ in our specifications.

Column "Within Section" in table 3 tabulates the parameter estimates. Consistent with depletion and well interference, the estimated well-order effect is negative, indicating that the output of subsequent wells in a section steadily declines as more wells are added, all else equal. The estimated learning signal-to-noise ratio κ of approximately 4 implies that the average output surprise from 4 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.

Our estimates allow us to decompose three mechanisms that account for how site selection affected wells' production: selection on observed geologic quality X_i (OOIP in the Bakken, OGIP in the Haynesille); selection on ex-ante unobserved geologic quality η_i (as captured by the control function $H(\hat{V}_i)$); and selection on firms' updated beliefs about geologic quality θ_i that are driven by learning from previous wells' output in the same section (captured by the κ-weighted average of previous wells' output surprises). Figure 12 plots the contributions of these mechanisms.[18] In North Dakota, the observed 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 lower unobserved quality persisted after 2012. Over time, learning from the outcomes of previously drilled wells eventually led firms to select better sites: relative to the pre-2012 average, learning added 0.05 log points to well

[18]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 11: Effect of the Realized First Well Production Outcomes

(a) Effect on the Probability of Drilling a Second Well

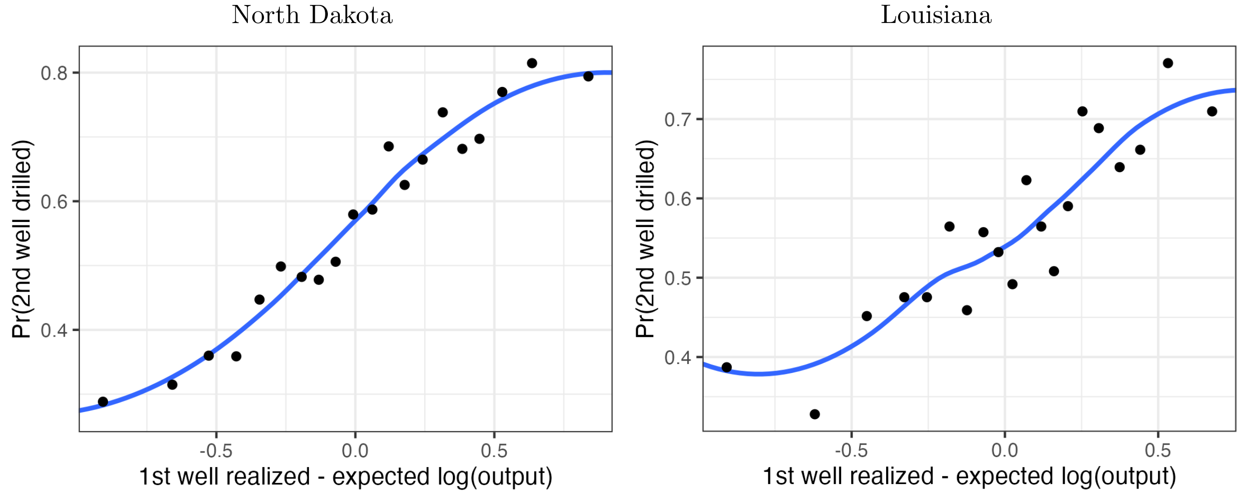

(b) Effect on the Realized Second Well Production Outcomes

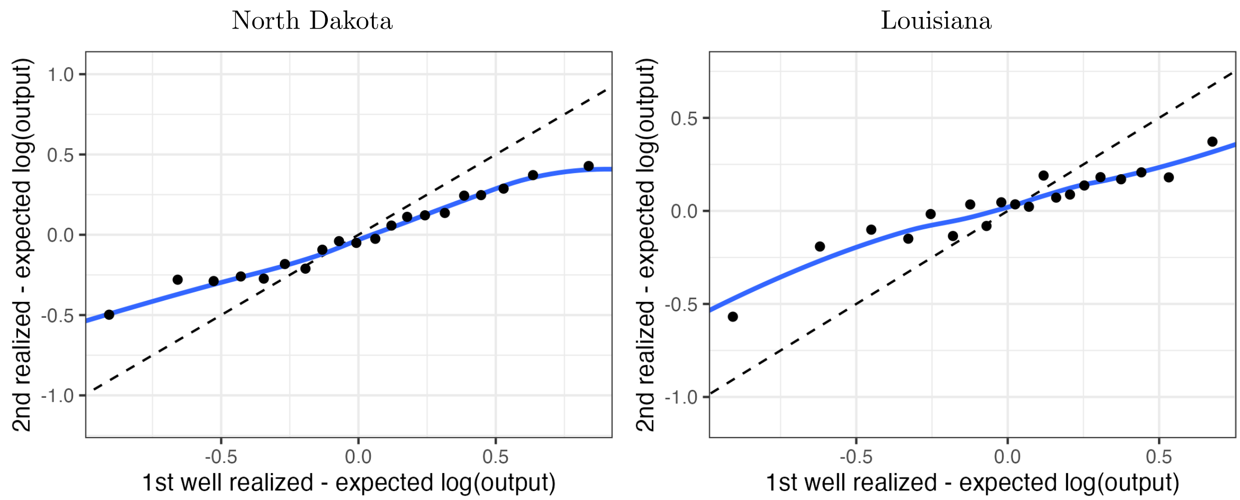

Notes: The filled circles display a binned scatter plot, constructed by dividing the sample into vigintiles (20 equal-sized bins) based on the x-axis variable. The solid lines represent LOESS curves fitted to the underlying, unbinned observation-level data. The dashed lines in Panel (b) indicate the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

production on average by 2016, rising further to 0.12 log points by 2017. In Louisiana, unobserved quality followed a similar pattern as North Dakota, sitting 0.04 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 after 2010, driving a 0.06 log point increase in production during 2013 to 2015.

We next move on to our full, across-section learning model. For *a priori* evidence that this

Figure 12: Decomposition of the Site Selection Mechanisms: Within-Section Learning Model

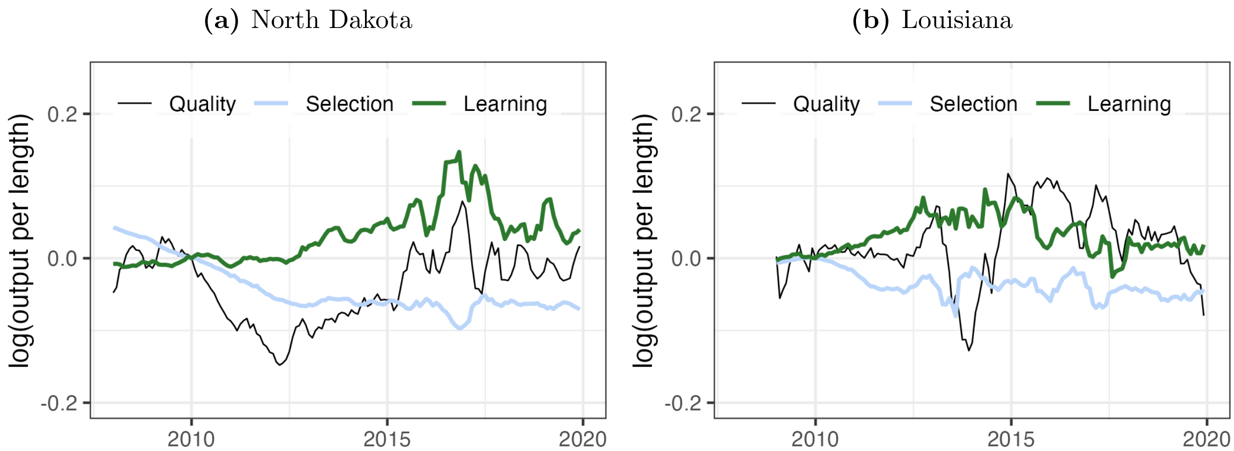

Notes: The dependent variable in each regression is the log output per lateral length. Each component plots centered five-month moving averages, calculated by averaging components across all wells drilled within each five-month window. The plotted range for North Dakota is from the beginning of 2008, and for Louisiana, from the beginning of 2009, reflecting that most wells are drilled after these periods. The log output per lateral length values are normalized to zero at the beginning of 2010.

learning is likely to be important, figure 13 shows that wells' output surprises are correlated across nearby sections. For wells drilled within six months of one another, the correlation of their production is highest within the same section and declines steadily with distance, nearing zero by a distance of six miles. These correlations motivate our spatial learning model: firms update their beliefs about geologic quality using wells drilled in nearby sections, with the strength of the updating decaying with distance up to a distance of six miles, beyond which we assume that the correlation is zero.

We present our estimates from the model that accounts for cross-section learning in column "Across Section" in table 3. The estimated learning signal-to-noise ratio κ is 8 for North Dakota and 6 for Louisiana, larger than the value of 4 estimated from the within-section model. Still, the larger number of wells that firms learn from in this model compensates for the lower signal precision, even though nearby wells' signals are discounted because their latent quality is imperfectly correlated with that of the section in question. The estimated latent quality correlation decay parameter γ_ρ implies latent quality correlations of 0.73, 0.53, 0.38, and 0.15 for sections 1, 2, 3, and 6 miles away in North Dakota, respectively, and correlations of 0.39, 0.15, 0.06, and < 0.01 in Louisiana. These estimates indicate meaningful across-section learning at close distances in both states, but with faster distance decay in Louisiana.

Figure 14 presents the same decomposition exercise as in the within-section model in fig-

Figure 13: Spatiotemporal Correlation of Output Surprises

(a) North Dakota **(b)** Louisiana

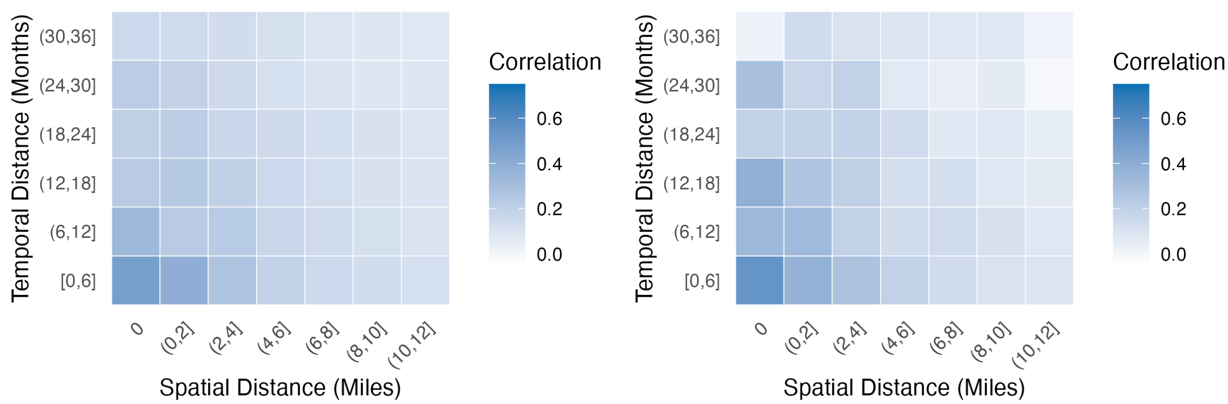

Notes: The figure plots pairwise correlations of output surprises for all unordered pairs of distinct wells within each spatial-by-temporal distance bin. Output surprises are defined as realized log production minus expected production from the estimated within-section learning model (including the control function term but without the learning term). Spatial distance is measured in miles between section centroids, and temporal distance is measured in months between drilling dates.

ure 12, now allowing the learning component to incorporate previous wells' output surprises from nearby sections. Compared to the within-section model, learning from previous wells is amplified and selection on unobserved ex-ante quality is attenuated, while the observed quality component changes little. The contribution of spatial learning to increased output is especially pronounced in North Dakota. We find that learning contributes an increase of 0.41 log points to production from 2010 to 2017, relative to an increase of 0.12 log points in the within-section model. The contribution of selection on ex-ante unobserved quality η is muted to nearly zero.

Figure 15 then plots the sum of the contributions of the three site-selection mechanisms, along with depletion effects $\lambda(j)$ and the contributions of technology and firms' operational decisions ($\omega(t)$). In North Dakota, the aggregate effect of the site-selection mechanisms shows a decline of 0.16 log points through 2012, as firms drilled initial wells in sections that were decreasing in quality. This result is consistent with the idea that firms wanted to drill at least one well in most sections prior to lease expiration in order to hold acreage and preserve the option for future drilling (Herrnstadt et al., 2024), and that they prioritized drilling higher-quality sections first. The decline in site quality was then followed by a bounce back, reaching an overall increase of 0.40 log points by 2017 relative to the pre-2010 period, once firms started to predominantly drill additional wells into previously exploited sections, focusing on higher-quality sections as they learned from previous wells' outcomes. In Louisiana, aggregate site-selection mechanisms had a smaller net effect on production,

Figure 14: Decomposition of the Site Selection Mechanisms: Across-Section Learning Model

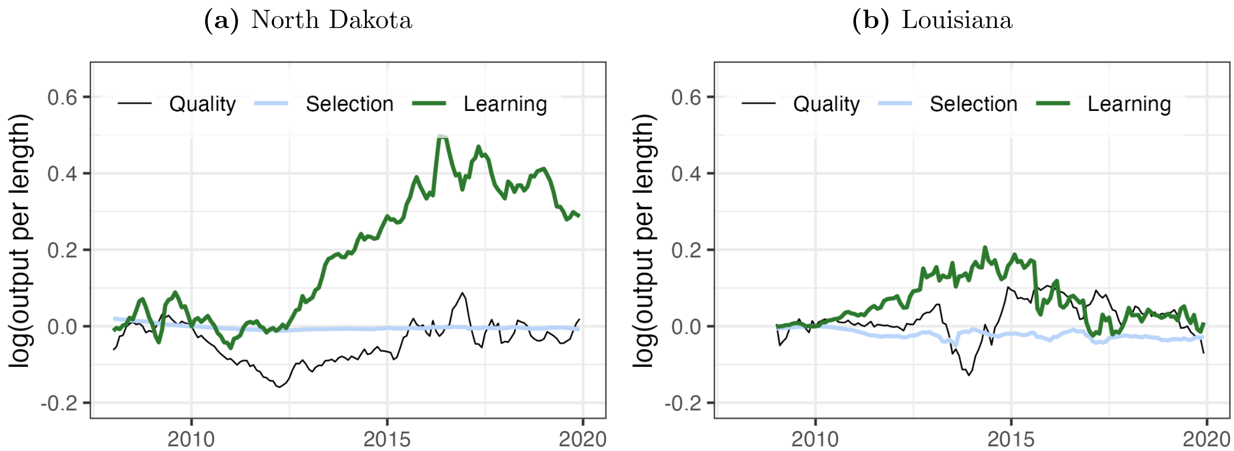

Notes: The dependent variable in each regression is the log output per lateral length. Each component plots centered five-month moving averages, calculated by averaging components across all wells drilled within each five-month window. The plotted range for North Dakota is from the beginning of 2008, and for Louisiana, from the beginning of 2009, reflecting that most wells are drilled after these periods. The log output per lateral length values are normalized to zero at the beginning of 2010. Note the change in y-axis scale relative to figure 12.

increasing output by 0.12 log points on average during 2013–2016 relative to the pre-2011 period, as selection on unobserved quality η offset some of the gains from learning.

Figure 15 also shows that well interference and depletion effects began to occur in North Dakota after 2012, amounting to a penalty of around -0.20 log points. These effects also occurred in Louisiana, imposing a decline amounting to -0.12 log points.

Notably, although site selection and depletion effects account for economically meaningful changes in output per well during our sample period, the overall decomposition still leaves the majority of the output growth captured by $\omega(t)$, which we attribute to changes in technology and firms' operational decisions. In North Dakota, $\omega(t)$ increases by 0.4 log points from the early period to 2011, stays relatively flat until 2014, and rapidly grows an additional 0.6 log points through 2018. In Louisiana, $\omega(t)$ follows a similar trajectory, rapidly increasing by 0.5 log points from the early period to 2010, staying relatively flat until 2014, and growing an additional 0.5 log points from 2014 to 2019.

Table 4 shows that our estimated average annual growth in $\omega(t)$ from the full model is 0.083 log points in North Dakota and 0.079 log points in Louisiana. If we allow for only within-section learning, the estimated average annual growth in $\omega(t)$ for North Dakota rises to 0.096, reflecting the reduction in the contribution of spatial learning to site selection. The estimate for Louisiana is substantively unchanged.

Table 4 also indicates how our estimates of the growth in $\omega(t)$ are affected when we do

Figure 15: Overall Output per Lateral Length Decomposition

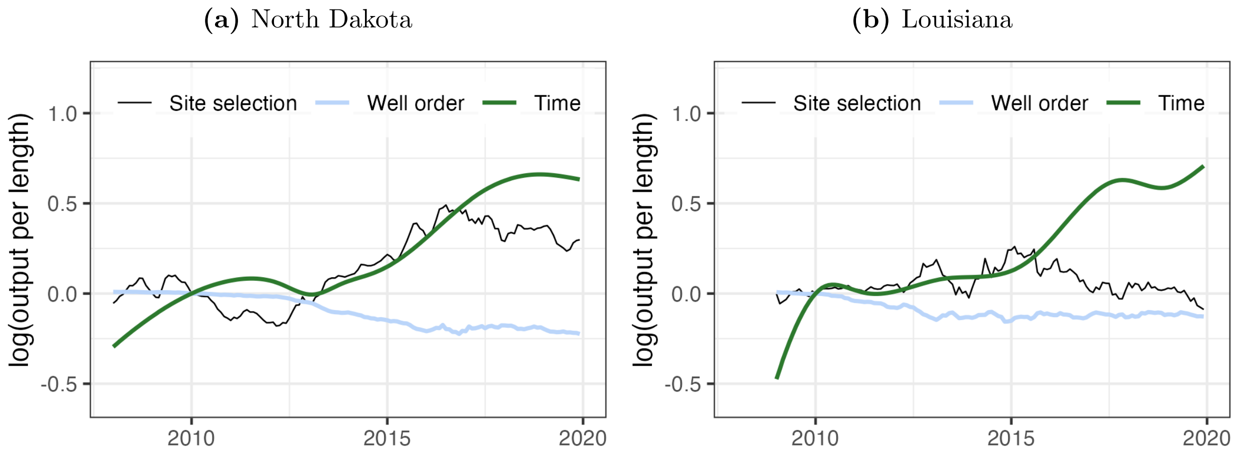

Notes: The lines for site selection and well order (depletion and well interference) are centered five-month moving averages, calculated by averaging components across all wells drilled within each five-month window. Site selection includes selection on observable X_i, selection on ex-ante unobserved quality η_i, and spatial learning. The time curve is evaluated directly from the estimated time-spline basis $\omega(t)$. 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 lateral length values are normalized to zero at the beginning of 2010.

not account for learning or the well order effects $\lambda(j)$. Focusing first on the case in which the learning model only allows for within-section learning, dropping the learning terms from the output equation does not significantly change the estimated growth in $\omega(t)$, despite the meaningful contributions of learning to site selection shown in figure 12. We obtain this result because the well order effects $\lambda(j)$ absorb some of the learning effect when learning is not explicitly accounted for in the model: they capture average unobserved quality differences across sections with different numbers of wells, and higher-quality sections tend to have more wells. This positive selection into later well orders offsets the negative depletion and interference effects that $\lambda(j)$ is meant to capture. As a result, the estimated $\lambda(j)$ function no longer has a clean economic interpretation, but it captures the net of learning and depletion effects well enough that the estimated growth rate of $\omega(t)$ is not substantially affected by the omission of learning from the model. If we exclude both learning and well order effects from the model, the estimated growth rate of $\omega(t)$ is reduced, as expected. Finally, in North Dakota we find a difference in the estimated growth rate between the spatial learning model versus the no learning model, since in this case the well order effects are not sufficiently correlated with production outcomes from nearby sections.

5 What explains the large increase in output per well?

Oil and gas drillers engaged in meaningful site selection, varying in both direction and magnitude over the coarse of the shale boom. After accounting for that selection, we find that firm's operational choices explain large increases in output per well over time. 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 first 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. Previous research has attributed large productivity gains in telecommunications (Olley and Pakes, 1996) and steel-making (Collard-Wexler and De Loecker, 2015) to a reallocation of production away from unproductive firms and towards more productive firms over time. To understand the potential role of firm-level selection in our setting, we re-estimate our main specification while including firm fixed effects in both the drilling hazard model and the output equation. These fixed effects control for time-invariant, firm-specific drilling and operational advantages.[19,20] 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 16 plots the contribution of firm FEs in the output per well decomposition. In both states, the estimated firm FEs exhibit a slight negative trend, the post-2014 average dropping by 0.02 log points relative to the pre-2011 average. Including firm FEs reduces the estimated average annual growth in $\omega(t)$ in North Dakota from 0.083 to 0.075 log points (Table 4), suggesting that changes in the composition of drilling firms explain a small part of the increase in output per well. Including firm FEs also attenuates the contribution of learning, which may result from firm FEs explaining part of the latent quality correlation in the model without firm FEs. In Louisiana, by contrast, the output trend estimates do not meaningfully depend on whether firm FEs are included or not, and the effects of the site selection mechanisms remain largely unchanged relative to the model without firm FEs. Thus, improved selection into more productive firms explains at most a limited share of the large increase in output per well over time. Instead, the gains are primarily attributable to operational choices *within* firms. This within-firm production growth contrasts with

[19]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.

[20]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 16: Output per Lateral Length Decomposition of the Firm Fixed Effects

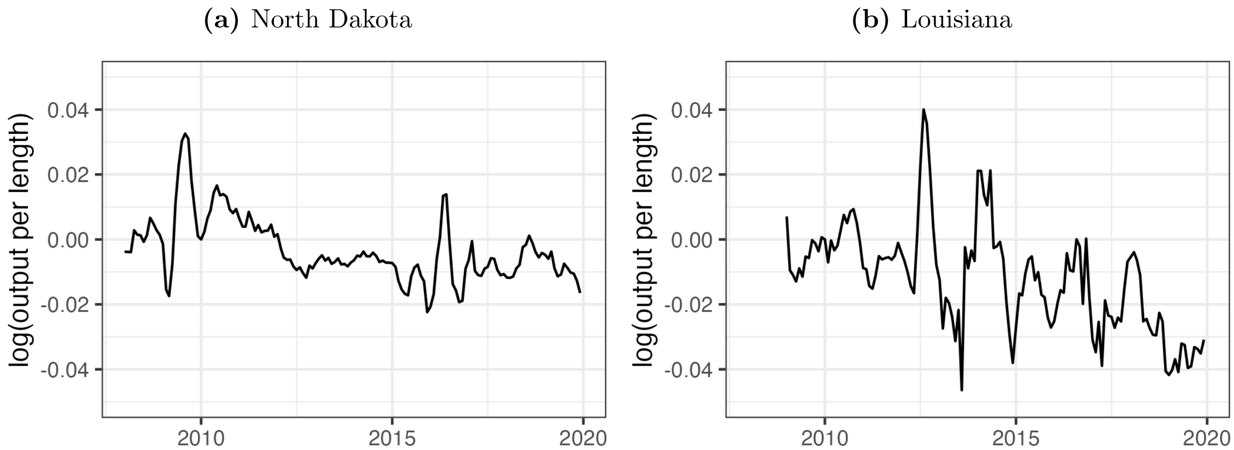

Notes: The plot uses centered five-month moving averages, calculated by averaging firm fixed effects across all wells drilled within each five-month window. The plotted range for North Dakota is from the beginning of 2008, and for Louisiana, from the beginning of 2009, reflecting that most wells are drilled after these periods. The log output per lateral length values are normalized to zero at the beginning of 2010.

consolidation or reallocation-driven aggregate productivity improvements documented in other settings.

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, beyond length. 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 inputs per foot drilled (figure 17). 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' 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.[21] 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 17: Proppant and Water Use Trends

(a) North Dakota **(b)** Louisiana

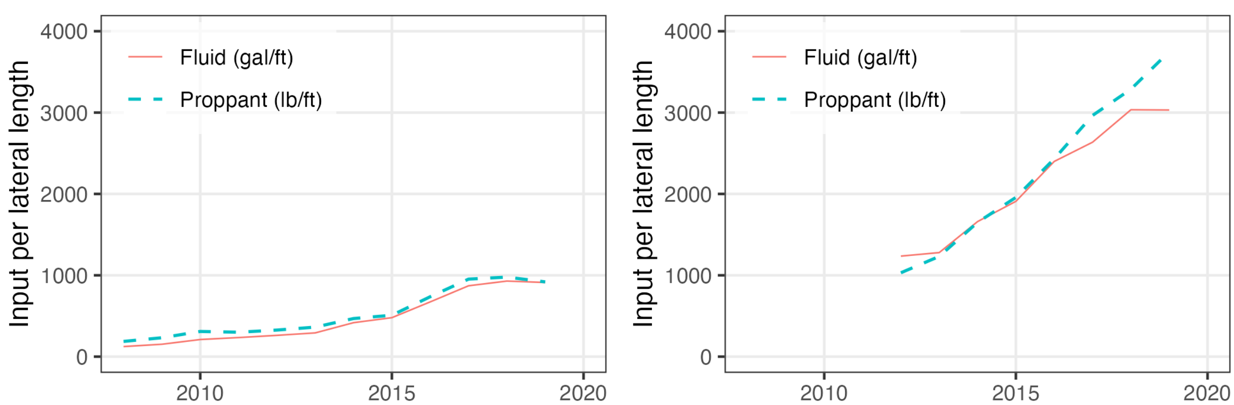

Notes: Figure shows the average annual input use calculated from Novilabs data. The plotted range for North Dakota is from 2008 to 2019. For Louisiana, the range is from 2012 to 2019, since fewer than 10 wells were recorded each year before 2012.

5.3 Changes in costs and input prices

Why did input use increase so much? Input use may have increased simply because the cost of inputs decreased. This would be a rational response on the part of firms, but would not indicate increased productivity. Another possibility is that firms were initially unaware of the

[21] 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 18: Rig rental day rates

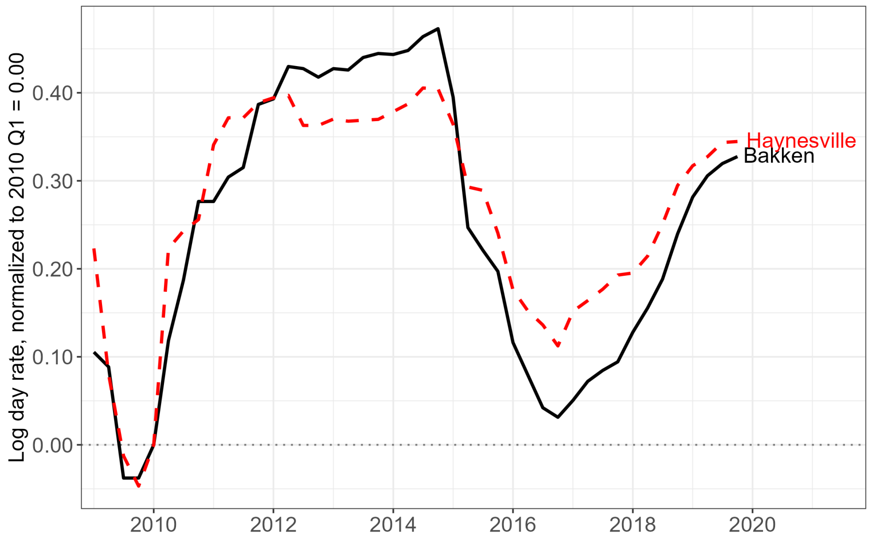

Notes: Figure shows quarterly day rate data from Enverus. We use Enverus's Rocky Mountain region for the Bakken, and "ArkLATex" for the Haynesville. Data correspond to depth ratings of 10,000–12,999 feet before 2018, and horsepower class 1000–1499hp from 2018 onward. Log day rates are normalized to zero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

large marginal product of key inputs, and learned over time that using more inputs like sand and water increased output and profits, as per the mechanism described in Covert (2015). Finally, firms could have adopted management practices or technologies that increased output, either via a standard Hicks-neutral TFP shifter or in a factor-augmenting way. For instance, adoption of equipment enabling multi-stage fracking would be complementary to increased use of water and sand. 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, increase 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 18. 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.[22]

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 (in logs) 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.[23]

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.[24] 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 19. 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. During this period, the average well lateral length increased by 30–80%, so that the cost per unit length fell during time in which output per unit lateral length was increasing.

To summarize, sand and water inputs per unit length more than doubled between 2010 and 2019. Input prices fell by 20–40% during 2015–2016, but, oil and gas prices more than halved during 2014–2016, suggesting this input price decline alone likely does not explain why firms increased observable inputs so much during this time period. When we look at a subset

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

[23] 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.

[24] 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 19: Reported drilling and completion costs per well

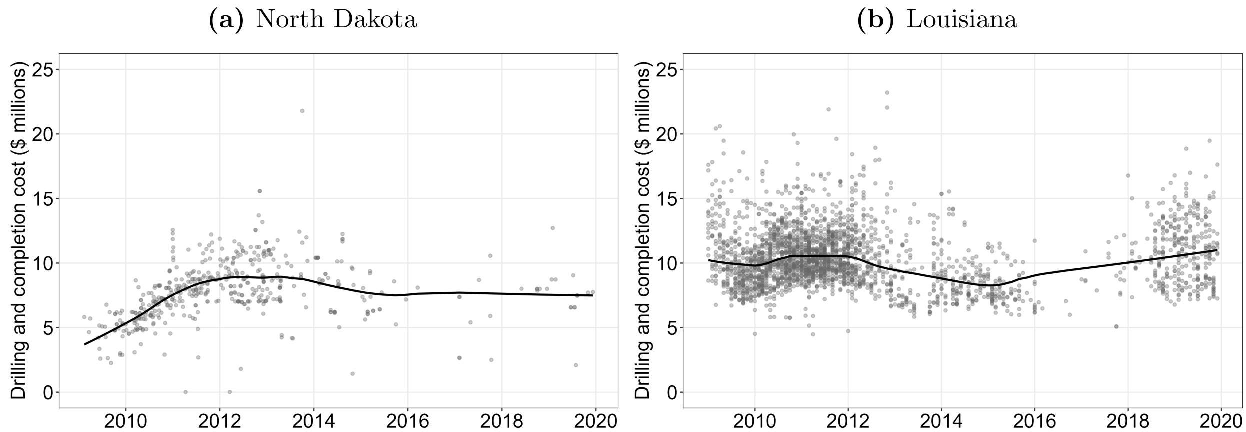

Notes: Figures show reported drilling and completion costs for a subsample of wells for which data are available. Each point is an individual well, and the solid line is a LOWESS curve fitted to the underlying data. Cost data are available and plotted for 509 wells in North Dakota and 2697 wells in Louisiana.

of available well costs, total costs are flat, while well lateral length increased by 30–80%. 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 they adopted new unobserved technologies that were complementary to input use. 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, so that sites drilled later tend to be lower quality. Then once many sites have been drilled at least once, firms preferentially drill sites that had better production outcomes, leading to an increase in site quality over time.

This learning effect eventually outweighs the first-well selection effect. At the same time, the repeated drilling of wells into the same sites adversely affects' wells' output through depletion and well interference effects. Overall, however, these mechanisms fall far short of explaining the large observed growth in output per well during the shale boom. Instead, evidence from output growth, input use, and completion costs implies that firms have made substantial improvements in the productivity with which they bring together inputs to produce oil and gas. This 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 a future clean 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, conditional on site, 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.[25] 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 industry, our approach for decomposing output growth into site selection versus technology and 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. Beyond energy, our approach could be applied

[25] 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).

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 the 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.

Ackerberg, Daniel A., Kevin Caves, and Garth Frazer, "Identification Properties of Recent Production Function Estimators," *Econometrica*, 2015, *83* (6), 2411–2451.

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).

Heckman, James J., “Sample Selection Bias as a Specification Error,” *Econometrica*, 1979, *47* (1), 153–161.

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.

Online appendix for "Investment, Productivity, and Selection in the U.S. Shale Boom"

A Additional model details

A.1 Un-modeled input use

This sub-section formalizes how changes in input use are incorporated into our estimates of the output equation (1), as discussed at the end of section 3.1 in the main text.

Suppose we have a single scalar input k_i, and suppose that the effect of k_i on log output is linear (as is commonly assumed in the production function literature). In expectation (building off of equation (5) in the main text), we can write the output equation that accounts for k_i as equation (A.1) below, re-labelling the original parameters with tildes to distinguish them from the parameters in the original output equation (5).

$$E[y_{ij}|t_j(i), X_i, Z_i, \mathcal{Y}_{ij}] = \alpha k_{ij} + \tilde{\omega}(t_j(i)) + X_i'\tilde{\beta} + \tilde{\lambda}_j + E[\tilde{\eta}_i|V_i] + E[\tilde{\theta}_i|\mathcal{E}_{ij}]. \tag{A.1}$$

We write k_{ij} as a function of time and of observed and unobserved section-level factors that are known to the firm:

$$k_{ij} = h(t_j(i)) + X_i'\pi_\beta + \pi_j + \pi_E(\tilde{\eta}_i + E[\tilde{\theta}_i|\mathcal{E}_{ij}]) + \nu_i, \tag{A.2}$$

where $h(t_j(i))$ is a flexible function of time that captures how time varying factors, including the technological progress in $\omega(t)$ but also other factors such as input and output prices, affect input use. The term ν_i represents unobserved section-specific factors that might affect input use (e.g. land constraints on lateral length, available road infrastructure, or environmental constraints) but do not directly affect output. Note that in this model ν_i can affect drilling timing (i.e., be correlated with V_i), since ν_i affects input use and thus output in a way that is forseeable by the firm.

Now re-write the expected output equation, substituting in for k_{ij}:

$$\begin{aligned} E[y_{ij}|t_j(i), X_i, Z_i, \mathcal{Y}_{ij}] = {} & \alpha(h(t_j(i)) + X_i'\pi_\beta + \pi_j + \pi_E(E[\tilde{\eta}_i|V_i] + E[\tilde{\theta}_i|\mathcal{E}_{ij}]) + E[\nu_i|V_i]) \\ & + \tilde{\omega}(t_j(i)) + X_i'\tilde{\beta} + \tilde{\lambda}_j + E[\tilde{\eta}_i|V_i] + E[\tilde{\theta}_i|\mathcal{E}_{ij}]. \end{aligned} \tag{A.3}$$

Unpacking the above yields:

$$E[y_{ij}|t_j(i), X_i, Z_i, \mathcal{Y}_{ij}] = (\tilde{\omega}(t_j(i)) + \alpha h(t_j(i))) + X_i'(\tilde{\beta} + \alpha\pi_\beta) + (\tilde{\lambda}_j + \alpha\pi_j)$$
$$+ E[(1 + \alpha\pi_E)\tilde{\eta}_i + \alpha\nu_i|V_i] + (1 + \alpha\pi_E)E[\tilde{\theta}_i|\mathcal{E}_{ij}]. \tag{A.4}$$

The above equation is what we are effectively estimating when we do not include inputs in the output equation. Thus, the parameters in the output equation (5) in the main text map to the parameters in the above equation as follows:

$$\omega(t) = \tilde{\omega}(t) + \alpha h(t), \tag{A.5}$$
$$\beta = \tilde{\beta} + \alpha\pi_\beta, \tag{A.6}$$
$$\lambda(j) = \tilde{\lambda}_j + \alpha\pi_j, \tag{A.7}$$
$$E[\eta_i|V_i] = E[(1 + \alpha\pi_E)\tilde{\eta}_i + \alpha\nu_i|V_i], \tag{A.8}$$
$$E[\theta_i|\mathcal{E}_{ij}] = (1 + \alpha\pi_E)E[\tilde{\theta}_i|\mathcal{E}_{ij}]. \tag{A.9}$$

A.2 Bayesian learning model

This sub-section formalizes the Bayesian learning model introduced in section 3.4, derives the posterior mean in equation (7) therein, and details the GMM estimation procedure. The key idea is that outputs from nearby sections can be informative about section i's latent quality θ_i, but that their informativeness declines with distance and is contaminated with the section- and well-specific noises in well output.

Let $\mathcal{E}_{ij}$ denote (a vector of) the history of output surprises (relative to the firm's ex ante expectation before drilling the first well in section i', accounting for depletion effects) available before drilling the jth well in section i:

$$\mathcal{E}_{ij} = \{e_{i'j'} : i' \in \{i, \text{ nearby } i\},\ t_{j'}(i') < t_j(i)\},$$

where

$$e_{i'j'} = y_{i'j'} - \omega(t_{j'}(i')) - X_{i'}'\beta - E[\eta_{i'}|V_{i'}] - \lambda_{j'}.$$

If θ_i and $\mathcal{E}_{ij}$ are jointly normal with mean zero, as

$$\begin{pmatrix} \theta_i \\ \mathcal{E}_{ij} \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_p^2 & \Sigma_{\theta_i,\mathcal{E}_{ij}} \\ \Sigma_{\theta_i,\mathcal{E}_{ij}}' & \Sigma_{\mathcal{E}_{ij},\mathcal{E}_{ij}} \end{pmatrix}\right),$$

then the conditional mean of the joint normal distribution implies that the Bayesian posterior mean is

$$E[\theta_i|\mathcal{E}_{ij}] = \Sigma_{\theta_i,\mathcal{E}_{ij}} \Sigma^{-1}_{\mathcal{E}_{ij},\mathcal{E}_{ij}} \mathcal{E}_{ij}.$$

This expression allows for a fully general covariance structure across latent section qualities and well-level signals. However, estimating it directly would require specifying and inverting a different, potentially high-dimensional covariance matrix for each well's history of output surprises. We therefore impose additional structure that preserves the core Bayesian intuition while yielding a closed-form section-level weighting formula.

Before drilling the first well in section i, the firm has a normal prior about the section's true quality, $\theta_i \sim \mathcal{N}(0, \sigma_p^2)$. We assume that the latent quality of a nearby section i', with discrete distance $d(i,i') \in \{0, 1, \ldots, D\}$ from section i, can be written as

$$\theta_{i'} = \rho_{d(i,i')}\theta_i + \nu_{i'},$$

where $\nu_{i'} \sim \mathcal{N}(0, \sigma_p^2(1 - \rho^2_{d(i,i')}))$, $\nu_{i'}$ is independent of θ_i, and $\nu_{i'}$ is independent across i'. The variance of $\nu_{i'}$ is pinned down by the normalization that each section's latent quality has prior variance σ_p^2. An output surprise (or a signal) from prior well j' in section i' is

$$e_{i'j'} = \theta_{i'} + \varepsilon_{i'j'} = \rho_{d(i,i')}\theta_i + \nu_{i'} + \varepsilon_{i'j'}, \text{ where } \varepsilon_{i'j'} \text{ i.i.d.} \sim \mathcal{N}(0, \sigma_s^2).$$

Let $\mathcal{J}^{ij}_{i'}$ denote the set of prior well indices from section i' in $\mathcal{E}_{ij}$, $J^{ij}_{i'} = |\mathcal{J}^{ij}_{i'}|$ denote the number of those prior wells, and $\bar{e}^{ij}_{i'}$ be the average signal from those prior wells $j' \in \mathcal{J}^{ij}_{i'}$. Here, the superscript ij indexes the drilling well and the subscript i' indexes the section to learn from. Then,

$$\bar{e}^{ij}_{i'} = \rho_{d(i,i')}\theta_i + \nu_{i'} + \bar{\varepsilon}^{ij}_{i'}, \text{ where } \bar{\varepsilon}^{ij}_{i'} \sim \mathcal{N}(0, \sigma_s^2/J^{ij}_{i'}).$$

For $\rho_{d(i,i')} \neq 0$, dividing by $\rho_{d(i,i')}$ yields an unbiased section-level signal about θ_i:

$$\tilde{e}^{ij}_{i'} = \frac{\bar{e}^{ij}_{i'}}{\rho_{d(i,i')}} = \theta_i + \frac{\nu_{i'}}{\rho_{d(i,i')}} + \frac{\bar{\varepsilon}^{ij}_{i'}}{\rho_{d(i,i')}}.$$

The variance of this unbiased signal, conditional on θ_i, is

$$\text{Var}(\tilde{e}^{ij}_{i'}|\theta_i) = \sigma_p^2 \left(\frac{(1 - \rho^2_{d(i,i')}) + \kappa/J^{ij}_{i'}}{\rho^2_{d(i,i')}} \right),$$

where $\kappa = \sigma_s^2/\sigma_p^2$. Thus, the precision of $\tilde{e}_{i'}^{ij}$, normalized by the prior precision $1/\sigma_p^2$, is

$$\pi_{i'}^{ij} = \frac{\rho_{d(i,i')}^2 J_{i'}^{ij}}{(1-\rho_{d(i,i')}^2) J_{i'}^{ij} + \kappa}.$$

The unbiased signal from section i', $\tilde{e}_{i'}^{ij}$, becomes more precise as the number of prior wells $J_{i'}^{ij}$ increases, but with diminishing marginal returns. The precision also increases with the spatial latent quality correlation $\rho_{d(i,i')}$ since signals from section i' are more informative about θ_i. If $\rho_{d(i,i')} = 0$, which indicates that signals from section i' contain no information about θ_i, then those signals have zero precision and do not enter the posterior mean. For the drilling section i, $d(i,i) = 0$ and $\rho_0 = 1$, so $\pi_i^{ij} = J_i^{ij}/\kappa$, which coincides with the precision in the within-section learning model (see equation (6) in section 3.4).

Since these unbiased section-level signals are independent across sections conditional on θ_i, the posterior mean becomes the precision-weighted average:

$$E[\theta_i|\mathcal{E}_{ij}] = \sum_{i' \in \{i' : \rho_{d(i,i')} \neq 0\}} \underbrace{\frac{\pi_{i'}^{ij}}{1 + \sum_\ell \pi_\ell^{ij}}}_{\text{Signal precision relative to prior}} \times \underbrace{\tilde{e}_{i'}^{ij}}_{\text{Unbiased signal}} = \sum_{i' \in \{i' : \rho_{d(i,i')} \neq 0\}} \frac{\pi_{i'}^{ij}/\rho_{d(i,i')}}{1 + \sum_\ell \pi_\ell^{ij}} \cdot \bar{e}_{i'}^{ij},$$

which is equation (7) in section 3.4.

We estimate this learning model by GMM. In estimation, we use a spline for $\omega(t)$ and polynomials for $H(\hat{V}_i)$ (control function for $E[\eta_i|V_i]$) and $\lambda(j)$. Collect these spline and polynomial bases and X_i into vector $\hat{W}_{ij}$, and let γ be the corresponding vector of coefficients. Let $\hat{e}_{ij}$ denote the estimated output equation residuals (taken with respect to expected output prior to observing the first well's outcome):

$$\hat{e}_{ij} = y_{ij} - \omega(t_j(i)) - X_i'\beta - H(\hat{V}_i) - \lambda(j) = y_{ij} - \hat{W}_{ij}'\gamma.$$

We take the average of these residuals of all wells in section i' before drilling the jth well in section i to empirically construct the section-level signal $\bar{e}_{i'}^{ij}$:

$$\hat{\bar{e}}_{i'}^{ij} = \frac{1}{J_{i'}^{ij}} \sum_{j' \in \mathcal{J}_{i'}^{ij}} \hat{e}_{i'j'}.$$

We then have the following estimation equation for the jth well in section i (same as equation (8) in section 3.4):

$$y_{ij} = \hat{W}_{ij}'\gamma + \sum_{i' \in \{i' : \rho_{d(i,i')} \neq 0\}} \frac{\pi_{i'}^{ij}/\rho_{d(i,i')}}{1 + \sum_\ell \pi_\ell^{ij}} \cdot \hat{\hat{e}}_{i'}^{ij} + \varepsilon_{ij},$$

which gives us the structural residual of

$$\varepsilon_{ij} = y_{ij} - \hat{W}_{ij}'\gamma - \sum_{i' \in \{i' : \rho_{d(i,i')} \neq 0\}} \frac{\pi_{i'}^{ij}/\rho_{d(i,i')}}{1 + \sum_\ell \pi_\ell^{ij}} \cdot \hat{\hat{e}}_{i'}^{ij}.$$

We use the spline and polynomial bases and covariates X_i, collected in $\hat{W}_{ij}$, as instruments to identify γ (which includes β and the parameters for $\omega(t)$, $H(\hat{V}_i)$, and $\lambda(j)$). To identify the learning parameters κ and ρ_d for $d = 1, \ldots, D$, we construct within-section learning instruments, defined as

$$Z_{ij}^0 = \left(J_{ij}^0, \bar{y}_{ij}^0, \bar{\hat{W}}_{ij}^0 \right),$$

and nearby-section learning instruments, defined as

$$Z_{ij}^d = \left(J_{ij}^d, \bar{y}_{ij}^d - \bar{y}_{ij}^{D+1}, \bar{\hat{W}}_{ij}^d - \bar{\hat{W}}_{ij}^{D+1} \right) \text{ for } d = 1, \ldots, D,$$

where $J_{ij}^d = \sum_{i' : d(i,i') = d} J_{i'}^{ij}$ is the total number of prior wells in sections at distance d from section i, and

$$\bar{y}_{ij}^d = (J_{ij}^d)^{-1} \sum_{i' : d(i,i') = d} \sum_{j' \in \mathcal{J}_{i'}^{ij}} y_{i'j'} \text{ and } \bar{\hat{W}}_{ij}^d = (J_{ij}^d)^{-1} \sum_{i' : d(i,i') = d} \sum_{j' \in \mathcal{J}_{i'}^{ij}} \hat{W}_{i'j'}$$

are the averages of outcomes $y_{i'j'}$ and combined covariates $\hat{W}_{i'j'}$ across those prior wells. We define these averages as zero when $J_{ij}^d = 0$.

These distance-wise instruments parsimoniously capture the (exogenous) observables involved in the empirical posterior mean:

$$\sum_{i' \in \{i' : \rho_{d(i,i')} \neq 0\}} \frac{\pi_{i'}^{ij}/\rho_{d(i,i')}}{1 + \sum_\ell \pi_\ell^{ij}} \cdot \hat{\hat{e}}_{i'}^{ij} = \sum_{i' \in \{i' : \rho_{d(i,i')} \neq 0\}} \frac{\pi_{i'}^{ij}/\rho_{d(i,i')}}{1 + \sum_\ell \pi_\ell^{ij}} \cdot \frac{1}{J_{i'}^{ij}} \sum_{j' \in \mathcal{J}_{i'}^{ij}} (y_{i'j'} - \hat{W}_{i'j'}'\gamma).$$

Intuitively, within-section instruments Z_{ij}^0 identify κ, and nearby-section instruments Z_{ij}^d identify the spatial latent quality correlations ρ_d for $d = 1, \ldots, D$. Nearby-section instruments use contrasts relative to a farther baseline distance $D+1$, assuming that $\rho_{D+1} = 0$ or

$D+1$ is sufficiently distant that it contains no information about the drilling section. In the estimation, we use $D = 3$: $d = 1$ for sections within $(0, 2]$ miles, $d = 2$ for sections within $(2, 4]$ miles, and $d = 3$ for sections within $(4, 6]$ miles. Sections in $(6, 8]$ miles compose the far-distance baseline $d = 4$.

The resulting GMM moments can be summarized as

$$E[(\hat{W}_{ij}, Z^0_{ij}, Z^1_{ij}, \ldots, Z^D_{ij}) \cdot \varepsilon_{ij}] = 0.$$

The model nests the within-section learning model as a special case when the latent quality correlations with nearby sections are set to zero, i.e., $\rho_d = 0$ for all $d = 1, \ldots, D$.