Learning Where to Drill: Drilling Decisions and Geological Quality in the Haynesville Shale

See Latest version

Mark Agerton *

January 1, 2019

Abstract

We often attribute the increasing productivity of U.S. shale oil and gas wells to firms learning *how* to drill better. Firms may instead be changing *where* they drill based on the interaction of their beliefs about geology with other economic variables. To identify what firms believe and learn about geology, and how this affects affects average output over time, I estimate an internally consistent model of royalty-rates, drilling decisions, and production outcomes in Louisiana's Haynesville shale. I find that some but not all of the increase in average output per well is explained by the structure of mineral leases and firms learning about where to drill.

*This paper was previously circulated as "Drilling to Learn: Information, Royalty Rates and Drilling Decisions in the Haynesville Shale." Special thanks to Edward B. Poitevent II (Stone Pigman Walther Wittmann L.L.C.), Silas Martin (Drillinginfo), Cullen Amend (Encino Energy), Jim McBride (Opportune LLP), Gregory B. Upton (LSU), Ryan Kellogg, Thomas Covert (University of Chicago), Peter R. Hartley, Xun Tang, Kenneth B. Medlock III, Mahmoud El-Gamal, Nick Frazier (Rice), Jim Griffin and Steve Puller (Texas A&M), Jim Wilen and Bulat Gafarov (UC Davis), Jevgenijs Steinbuks (World Bank), Eric Lewis (DOJ).

1 Introduction

The US shale boom has been a convenient setting for economists to study how firms learn about a new production process. The primary outcome of interest for these papers is how output per well increases with industry or individual drilling experience. The narrative being told is that technological progress has brought us into an age of hydrocarbon abundance. The sites that firms choose drill, however, are not randomly chosen with respect to the quality of their underlying geology. Instead, geology interacts with a cornucopia of other factors that vary over time and space to determine the full economic benefit of drilling. When we fail to fully account for the process of site selection and the role of geology in it, we may mistake changes in where firms drill for improvements in how they drill. Because oil and gas are depletable natural resources, the distinction between whether we have been learning how to drill or changing where we drill has bite. Learning how to drill makes all locations produce more, expanding the capacity of the resource base. In contrast, concentrating our drilling on only the most prolific locations simply shifts production forward in time, accelerating the rate at which we deplete the resource. Should we confound the roles of how and where to drill, we run the risk of overstating our long-run resource supply.

Of the papers that examine well productivity in shale, Fitzgerald (2015) and Steck (2018) control for geology at the level of a coarse, six by six mile area, and Covert (2015) and Montgomery and O'Sullivan (2017) use sophisticated spatial econometrics to estimate its contribution to production. All four assume that well locations are chosen randomly. Geology, however, varies rapidly over space (Covert 2015; Montgomery and O'Sullivan 2017), and firms know a great deal more about it than do we econometricians, so geology represents a form of unobserved heterogeneity when we examine a sample of production data. It interacts with several other factors to determine the full economic payoff from drilling a location and, therefore, the probability that a location is drilled. These economic factors include price and cost, the terms of mineral leases, the royalty rates firms must pay mineral owners, and the value of acquiring new information about a location's underlying geology. To further complicate matters, royalty rates may be positively correlated with firms' beliefs about geological quality at the same time as they reduce firms' payoff to drilling. We have less than a decade of data on the U.S. shale boom, and over this short period, cumulative experience is correlated with structural shifts in prices, changes in the incentives to drill implied by mineral lease contracts, and improvement in firms' information about geology. It is very possible that some of the *apparent* increases in well productivity that we have documented are due in part to changes in *where* firms drill, not *how* they drill.

In contrast to the papers mentioned above, I ignore the role of improvements in *how* firms drill and concentrate exclusively on whether firms' selection of *where* to drill individual wells also explains rising aggregate output per well. My central challenge is that we cannot observe what firms believe about the quality of the locations they drill. To overcome this information deficit, I construct an internally consistent model that allows me to identify the distribution of unobserved heterogeneity (geology), how it affects firms' decisions, and what firms learn about geology over time. The model has three parts. First, a firm negotiates a royalty rate for a lease given its initial priors about the location's geological quality. Second, the firm optimally decides when to drill given its current information about geology, prices, costs, and the terms of its mineral lease. It must balance the benefits from today's sure production revenues plus the value of new information about geology with the costs of paying for a well and forgoing the option to drill later. This second component is the heart of the model and uses the Rust (1987) dynamic discrete choice framework. Third, observed production volumes depend on a location's geological quality.

The model maps unobserved heterogeneity in geology into observable outcomes in the following way. Conditional on exogenous variation in mineral owner characteristics, firms should be willing to pay high royalty rates. Then, conditional on exogenous variation in prices, costs, and the structure of mineral leases, firms accelerate initial drilling when they expect an area to be very productive. They also accelerate initial drilling if they know they can resolve a great deal of uncertainty about geology. Should firms learn that an area does in fact have very favorable geology, firms are likely to accelerate drilling additional wells and will drill more of them. These wells will produce observably a larger volume of hydrocarbons. I assume that all wells within the same square mile share the same geology. This implies I observe multiple measurements for how unobserved geological heterogeneity affects production volumes, and the restriction further assists with identification.

I turn to Louisiana's Haynesville shale to estimate the model because the state partitions space into uniformly sized, one square mile sections that organize firms' drilling decisions and mineral owners' property rights. Since each section can hold up to eight wells, I can see the wells firms choose *not* to drill in addition to the ones they do. Descriptive statistics from the Haynesville reveal two statistical regularities that are consistent with my model about selection. First, firms accelerate drilling in locations where they agree to pay higher royalty rates. Because high royalty rates disincentivize drilling, firms must believe these locations to be extra-profitable. Second, wells in locations with higher royalty rates produce more natural gas.

Estimation of the structural model reveals several important aspects of firm behavior. First, firms' initial priors about geology are informative but not perfect: the correlation of these with actual quality is 0.66. This implies that drilling an initial well provides additional informational value in addition to a financial payoff. Compared to a perfect information scenario, preliminary estimates of learning can rationalize an additional 15-20% increase in productivity per well over the 2008–2016 period. Second, the expiration of mineral leases provides a powerful incentive for firms to drill at least one well in lower-productivity locations. As firms complete drilling of these initial wells, simulations suggest that the model can rationalize increases in aggregate productivity up to approximately 30% over an eight year time span. After I correct for the effects of firms selecting drilling locations based on geological quality, the total contribution of an exogenous technology trend over the period 2008–2016 falls from 39 to 34%. While the difference in parameter estimates is not statistically significant, it is economically significant and emphasizes the importance of controlling for the role of how firms choose which wells to drill based on information that impacts output but cannot be directly unobserved by analysts.

2 Literature review

The oil and gas industry has long been a fruitful environment to study learning since we can observe many small investment decisions that firms make, as well as the outcomes of those investments. One strand of research uses the shale context to study how firms learn about a production process as they drill wells over time. Covert (2015), Kellogg (2011), and Steck (2018) all examine how firms learn from (or with) one another, and Fitzgerald (2015) and Seitlheko (2016) document that firms with more experience see better well performance. In each of these studies, learning implies either increases in output per well or lower drilling costs. Such gains from learning are presumed to be transferrable across drilling sites, ultimately increasing the productive capacity of the resource.

In contrast to the above papers, the model I construct and estimate assumes that there has been no technological improvement in natural gas extraction. Instead, I assume that changes in average well output over time are due to changes in the way firms select which grades of geology to drill. My estimates are best interpreted as a "worst-case" scenario for long-run supply in which technological progress can do nothing to reverse the effects of depletion, just as the aforementioned papers are "best-case" scenarios that assume the resource is not finite. Similar to this paper, Smith (2017) and Smith and Lee (2017) provide a simple method to correct for the way firms' selection and subsequent exhaustion of better deposits affects the elasticity of supply. However, these two papers are not able to control for the full range of economically important factors like mineral leases, nor are they able to explicitly link individual drilling decisions and production outcomes as I do.

One of the benefits of focusing on Louisiana's Haynesville shale is that advances in extraction technology are likely to have had less impact compared to other settings. Much of the investment activity in the Haynesville occurred during the early years of the shale boom before 2012. After this, activity shifted to oil-rich shale plays like North Dakota's Bakken Shale and Texas' Permian Basin. Much of the technological development in shale extraction had to do with adapting techniques that worked in gas plays to oil plays. Another important development in shale technology was the ability to drill longer wells. Since well length is limited by the institution governing mineral rights in Louisiana, this new technology and firms ability to choose well length has had less impact.

My paper is closer to another strand of research that holds the production process fixed and studies how firms use information in making their drilling decisions. A working paper by Levitt (2009) casts 38 years of firms' drilling decisions in the Canadian province of Alberta as a private, Bayesian learning process. In Levitt's model, firms update their priors about geological quality with each new well they drill. Because I focus on the Haynesville, I am able to exploit a richer dataset that provides more information to identify learning. Furthermore, my shorter time period and smaller spatial area reduce the scope for unobserved heterogeneity and model mispecification to drive my results.

Another group of papers by Hendricks and Kovenock (1989), Hendricks and Porter (1996), and Lin (2013) turn to offshore drilling study the effect of information spillovers on the timing of firms' decisions. They find that these spillovers induce delayed drilling. Intuitively, a firm wants to wait for its neighbors to drill first so that the firm can avoid drilling an unprofitable well. Fortunately, the issue of information spillovers is likely to be limited in my setting: firms will delay only the first of the eight possible wells, and the expiration of mineral leases limits the amount a firm can delay drilling. The potential bias introduced by ignoring information spillovers should cause my estimates to understate the true extent of learning and overstate the precision of firms' prior beliefs. The intuition for this is as follows. When firms expect to learn less from the first well, they delay drilling. We can rationalize such an empirical delay by understating the informational gains from drilling (i.e., learning). Such a bias will limit the degree to which the model predicts that learning where to drill contribtues to productivity growth. This bias will not extend to the parameters which govern long run depletion once firms know about the spatial distribution of geology.

The problem of when to drill (or decomission) a well is a dynamic, discrete choice problem that Kellogg (2014), Muehlenbachs (2015), and I all study. Like them, I use the machinery of the Rust (1987) model to estimate parameters that characterize firms' payoffs and beliefs. By adding information on royalty rates and drilling outcomes, I show how one can combine some of the additional data available in shale plays to incorporate a greater degree of unobserved heterogeneity than before.

Like Herrnstadt, Kellogg, and Lewis (2018), I find that the terms of mineral leases are major drivers of firms' drilling decisions in the Haynesville: the expiration of primary terms and lease extensions induce spikes in investment, and high royalty rates deter it. My paper differs in several important dimensions, however. First, my focus is on firms' drilling and learning problem, not mineral owners' asymmetric information problem. Second, I take a much less of a stand on the game being played between mineral owners and firms. Instead, I use reduced form methods that can capture the relationship betwteen observable royalty rates, mineral owner characteristics, and unobserved geology. Third, the theoretical model in Herrnstadt, Kellogg, and Lewis (2018) implies that higher royalty rates must correspond to higher levels of uncertainty around geological quality. Contrary to their theoretical result, I find that higher royalty rates are in fact associated with more productive locations, and I assume that the distribution of geological quality is the same across all locations.

3 Institutional details

Louisiana's Haynesville shale is an especially good place to study firms' investment decisions because we see the investments firms choose *not* to make, in addition to the ones that they do. The state regulator partitions the Haynesville into roughly one square mile (640 acre) blocks called *sections*. Each section requires around eight wells to fully exploit. The partition into sections is based on the Public Land Survey System (PLSS) grid created during the 19th century, determined long before firms' present-day decisions. Figure 1 shows this partition, as well as the formation's location in northwest Louisiana. When a a firm wants to drill and produce a well in a section, the state forms a *drilling unit* that coincides with the section. While only one firm is allowed to make decisions about a well (the *operator*), all parties with mineral interests in the unit must participate in the well, be they original mineral owners or firms who have leased the rights. Thus, these pre-defined, square-mile sections systematically partition the shale into discrete sets of investment opportunities with one decision-maker in each one. Because shale formations exhibit low permeability, hydrocarbons do not flow into wells from very far away. Low permeability implies that wells in one section do not drain hydrocarbons from a neighboring section. This limits the scope for common-pool externalities (such as one firm draining deposits under its neighbor) to affect drilling behavior.

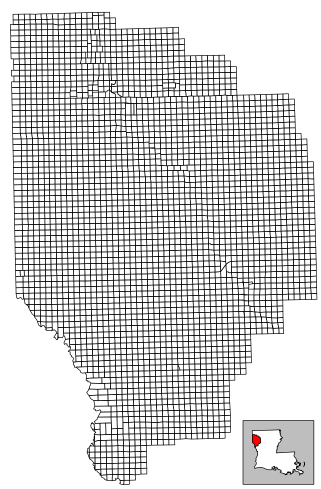

Figure 1: PLSS sections in Louisiana's Haynesville shale

Operators can only drill wells that originate on surface locations under which they hold mineral rights, and they normally attempt to lease the majority of a section before drilling.[1] Ownership of the mineral rights within

[1] Special thanks to Edward B. Poitevent II (Stone Pigman Walther Wittmann L.L.C.), Silas Martin (Drillinginfo), and Cullen Amend (Encino Energy) for assistance with the institutional background of mineral leasing.

a section is generally split among multiple private individuals; rights are normally not state-owned. Firms approach mineral owners, either directly or through third-party "landmen," and negotiate mineral leases with them. These bilateral contracts give the firm the right (but not obligation) to drill wells and sell the minerals produced produced. In exchange, the firm agrees to pay the mineral owner an up-front, cash payment, called the *bonus bid*, and a percentage of any revenue recieved from selling extracted minerals, called the *royalty rate*. Bonus bids are generally not reported, but public mineral lease records often specify the royalty rate. Rates in the Haynesville range from 12.5% to 25%, with more recent leases tending to be in the 20–25% range (see Figure 15 in the Appendix). While a higher royalty rate can raise the landowner's revenue, it also reduces the firm's incentive to drill.

Mineral lease contracts and a firm's right to drill an initial well expire after an initial *primary term*, usually three to five years. Around 80% of leases in my sample allow firms to extend the primary term in exchange for a cash payment. These *lease extensions*, also called "kickers," normally last two years. Should the firm drill and commence production within the primary term, the lease is considered to be *held by production*, and it enters into an indefinite *secondary term*. Since all mineral interests in a unit must participate in each well, all leases in the unit are held by production, even if they do not physically contain a well. The secondary term lasts as long as production continues in commercial quantities (Lane, Freund, and McNab 2015). During the secondary term, the firm retains the right to drill additional wells at any time. Such a contract structure implies that the economic payoff to drilling an initial well can be quite large, as the cost of drilling an initial wells provides financial payoff, information about geology, and an option to drill several more wells. In the sample of sections I study, only 22% see no drilling by October 2016 (see Table 6 in Appendix A). Lease expirations will serve to accelerate drilling compared to the case where firms own minerals outright as the opportunity cost of forgoing the option to drill again can be quite large.

4 Data

Firms make investment decisions at the level of a section, so I take a section as my unit of observation. There are three stages of activity on each section: mineral leasing, drilling, and production. In the following paragraphs, I briefly discuss how I assemble the data and define my sample. Details on each step of the data assembly process are available in Appendix A.

A shale play is defined by its geology, so I define the geographic extent of Louisiana's Haynesville shale using a University of Texas Bureau of Economic Geology study on the geological quality of the Haynesville (Browning et al. 2015; Gülen et al. 2015) plus a three mile buffer around it. The Louisiana Department of Natural Resources (DNR) provides GIS shapefiles of sections and drilling unit polygons, and I limit my attention to sections within the Haynesville. I use these polygons to partition space into sections $i = 1, \ldots, N$.

I use DNR GIS data on the locations of wells drilled in the Haynesville region, and I remove all conventional wells that are not shale-directed. These are mainly smaller wells drilled years prior to the shale boom. Such wells use a very different production process compared to shale wells, and they access much shallower layers of rock. I then use a combination of wells' spatial locations and names to match them to the sections they are associated with. I take well characteristics from the DNR's SONRIS database, and I merge the DNR well data with production data from commercial-provider Drillinginfo. I also merge in average monthly futures prices from Bloomberg, and PPI indices for drilling (PCU213111213111) and final demand less food and energy (WPSFD4131), which I take from FRED.

My goal is to focus on the sequence of drilling decisions that firms make *after* leasing, not the leasing process itself. There are two main characteristics of leases that affect firms' decisions: the start and expiration dates that constrain firms' ability to drill and the average royalty rate that the firm must pay mineral owners. The State of Louisiana requires that mineral leases be filed in the local parish courthouse. Drillinginfo, from which I obtain my leasing data, sends employees to parish courthouses to record the details of each lease signed in 2002 or later. The information recorded always includes the spatial location, date of the lease, the primary term, and whether an optional extension was specified. Royalty rates are included for leases, but not memorandums of leases.

I spatially merge leases with the sections they overlap. In what ends up as my final sample, 83% of leases are fully contained within a single section, and 12% span two sections (see lease-level summary statistics Table 9 in Appendix A). When a lease does overlap more than one section, I assume that drilling in one section cannot hold the share of the lease contained in another section. Each section usually contains many leases: in my final sample, the median and mean number of leases in a section are 10 and 18, respectively, and the maximum number is 355 (see section-level summary statistics Table 6 in Appendix A).

I next temporally merge primary terms with initial wells. I define extended primary terms by the start of each lease and its (possibly extended) expiration. Then I match the first shale well drilled in a section to the set of primary terms it overlaps and compute how much time remained in the primary term when the first well was drilled. In sections that see at least one well drilled, around 14% of mineral leases either expire before the first well or are signed after drilling starts. These lease/section pairs do not affect the sequence of drilling decisions we observe, so I drop them from the sample.

Wells that are drilled within a short time of one another are likely from the same investment decisions. Drilling a well tends to take from two to four weeks, and well completion takes additional time. Drilling multiple wells at once also helps lower costs since moving equipment is costly. To accommodate this behavior, I denote any well drilled within 8 weeks (less than 63 days) of another as belonging to the same drilling decision. For example, if 3 wells are drilled two weeks apart, I assume that the firm made one decision to drill three wells instead of three decisions to drill one well. Figure 20 in the Appendix shows the distribution of weeks since the previous well was drilled and where the 8-week cutoff lands. I then aggregate up time-varying variables like prices and the number of wells drilled to a quarterly frequency. Drilling decisions are not made very frequently, and this aggregation corre-

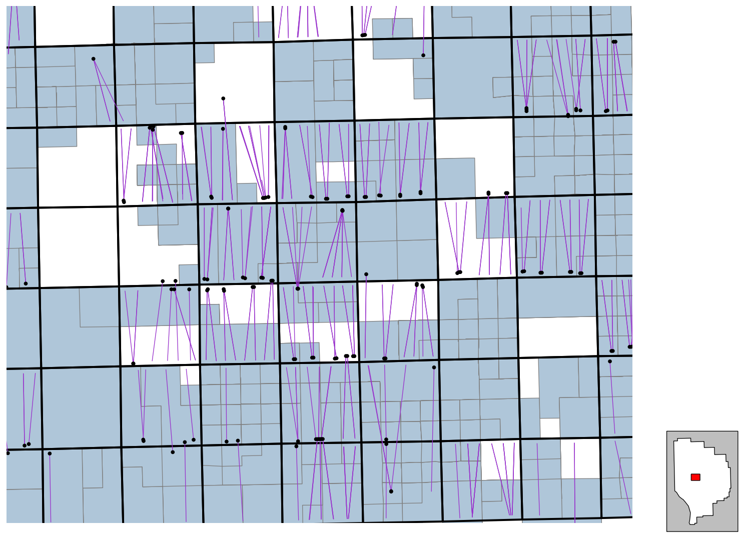

Figure 2: Wells, leases, and sections

sponds to a more appropriate frequency than weekly or monthly variation. It also raises the probability that a well is drilled in a given quarter, which helps in estimation.

The other key element of mineral lease contracts that affects firms' investment decisions is the weighted average royalty rate in a section, where the weights correspond to the percent of the section corresponding to each lease. The locations and sizes of leases are recorded at a coarse level, so the mineral lease polygons provided by Drillinginfo may *appear* to overlap or coincide, even when the actual leases do not. To avoid double-counting the area leased by firms, I weight the overlapping areas of each lease by one over the number of overlapping leases.[2] I assume that the contracts without royalty rates (usually Memos of leases) are drawn from the same distribution as those with royalty rates. Almost all of the royalty rates fall into one of six discrete categories: 12.5%, 16.67%, 18.75%, 20%, 22.5%, and 25%. I compute a weighted average royalty rate for each section and then map average royalty rates back to the nearest discrete one.

[2]Figure 14 in Appendix Ashows a simple example of how two partially overlapping leases would be handled.

Figure 2 shows a map of how these data fit together in a small area within the Haynesville. The squares with heavy, dark outlines are the PLSS sections. Within each, the faint blue shapes represent the outlines of mineral leases of varying sizes. These generally fall within section-boundaries. Wellheads (the surface location of the vertical wellbores) are marked by round dots, and these are connected via the purple line-segments to bottom-holes (the end of the horizontal portion of the well).

The next step in constructing my dataset is to merge in section-level covariates. The first is a measure of location-specific natural gas content from the Browning et al. (2015) and Gülen et al. (2015) study of geological quality in the Haynesville shale. Using geological data sampled throughout the shale, the authors estimate the spatial distribution of geological fundamentals like the thickness and total organic content. From these, they calculate a measure of estimated "original gas in place" (OGIP) in billion cubic feet per square mile for a grid of one-mile squares (see map Figure 16 in the Appendix). OGIP is calculated using geological fundamentals, so it is not endogenous to the history of firms' decisions as production data would be. Moreover, since OGIP is based on the sort of coarse geological information that firms should have access to, I will assume that the variable is in their information set before they start drilling.

The second set of covariates has to do with land use and surface characteristics for each section. I take section-specific averages over satellite-based land-cover data from the U.S. 2001 National Land Cover Database, the urban/rural classification from the 2010 U.S. Census, and the 2001–2006 average Census block-group characteristics from the American Community Survey (ACS). Figure 17 displays the satellite data on land imperviousness, along with the outlines of the state of Louisiana and the Haynesville shale (white), as well as the Census-defined urban areas (blue). I also use address information on mineral lease owners to compute the share of acreage in each section that is owned by out-of-state individuals.

From a universe of 3158 sections, I drop a little over half for one of nine different reasons: 754 for a single reason, and 990 for multiple reasons. This leaves 1414 sections in my final sample. The reasons for dropping sections relate to missing data, whether a section is in an urban area, and whether the section fits into a standard pattern of development, with leasing happening after 2003 and subsequent shale drilling starting during a primary term. Figure 18 and Table 5 in Appendix A provide more detail on why particular sections are dropped.

5 Descriptive evidence

Figure 3: Haynesville development over time

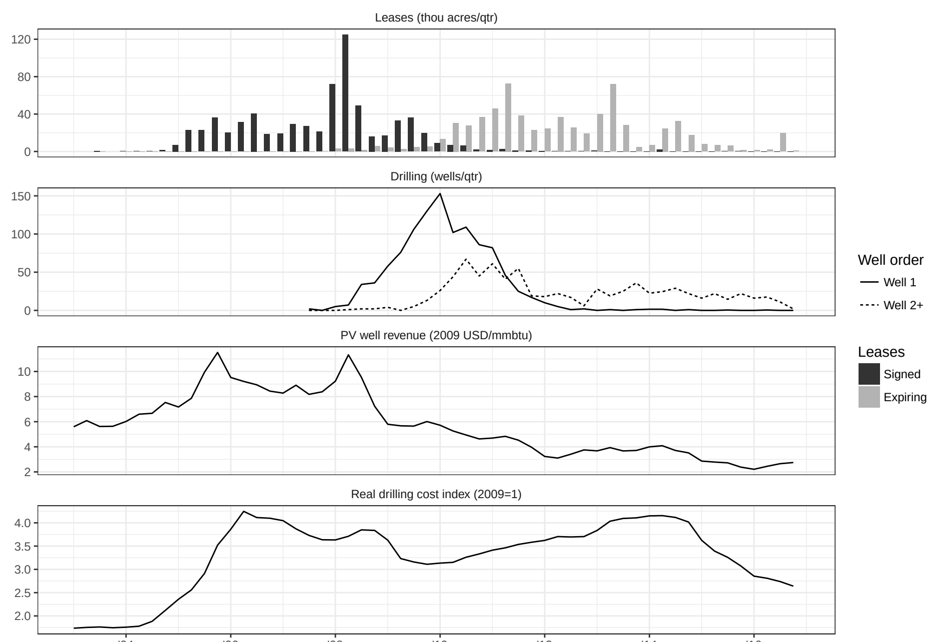

For many years, firms knew that gas deposits existed in the Haynesville shale formation but were not able profitably extract the gas. Then, in the early-to-mid 2000s, new technologies allowed firms to start producing gas from a similar, nearby formation, Texas' Barnett shale. Soon, firms' attention turned east towards the Haynesville, and by 2008, a "land-rush" (actually, a mineral rights rush) was on. The panes of Figure 3 plot the history of investment from 2003 to 2016. The top pane shows quarterly mineral leasing when leases expire. The frenzy of leasing in 2008 coincides with a peak in gas prices, which are shown in the bottom pane. The middle pane breaks out the number of wells drilled per month by whether a well is the first in its section, or whether it is drilled subsequently. As the plot shows, firms tend to delay drilling—both initial and subsequent—for a few years after leases are signed. By the time drilling picked up in 2009, prices were already falling. The rate at which firms drilled initial wells reached its zenith in the first quarter of 2010 before a raft of lease expirations. Drilling of the subsequent wells 2–8 in a section continued at a much slower pace. The fact that firms would drill initial wells as prices were falling has four explanations. First, the incentive to learn about geology may have pushed firms to drill initial wells despite lower prices. Second, the "use-it-or-lose-it" deadlines imposed by mineral lease expirations would have motivated them to drill. Third, as shown in the bottom pane of Figure 3 shows, costs were lower during the 2009–2011 period. Fourth, that firms did not start drilling until prices had fallen far from their 2008 peak suggests that firms may needed time to turn their focus from leasing to drilling. They may have faced high internal adjustment costs or high external drilling costs associated with a limited supply of drilling services.

Figure 4: Quarterly drilling hazard by well-order for 3 year leases

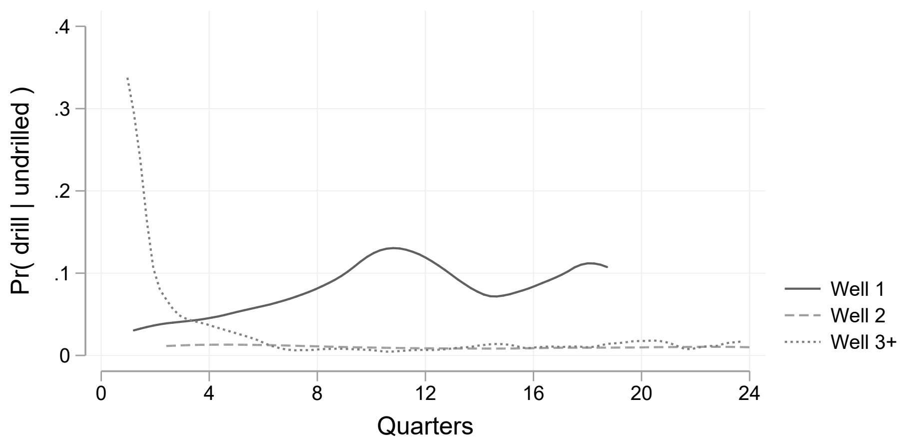

Standard real options theory suggests that the option value of waiting should fall as the expiration of a mineral lease approaches. As a result, the probability of drilling an initial well should increase as leases expire, and this is what we see empirically. Figure 4 shows an estimate of the hazard rate for drilling on three year leases, separated by whether the well is the first, second, or subsequent one in a section.[3] The horizontal axis starts at zero weeks when a lease is signed. The estimated hazard for initial wells peaks just before three years (12 quarters), the prevailing lease length, and again just before five years (20 quarters), when many lease extensions expire. Because of the empirical importance of these expiration dates, I feature them prominently in the model of drilling that I construct.

The level and shape of the hazard rates for well one contrasts sharply with that of well two. The drilling rate for Well 1 increases dramatically around expiration dates, and the drilling rate is also uniformly much higher compared to well two. This is consistent with three explanations. First, an option that expires in a finite time has less value than one that does not expire. This should make firms more willing to drill well one than subsequent wells. Second, the new information about geological quality that well one provides should make firms more likely to drill it. Third, prices were generally higher when firms were drilling their initial wells. Distinguishing between these three factors requires us to model the structure of firms' problems. The hazard rate for wells three to eight is different yet again from the hazard rates of wells one and two. The initial maximum and subsequent sharp decline in the well three hazard suggests if a firm does go ahead and drill well two, it is highly likely to drill additional wells.[4] Moreover, these are generally drilled one right after another. Such behavior could be consistent with fixed costs that are associated with starting up drilling in a location. Firms may be able to lower average costs by drilling multiple wells at once.

[3] The unit of observation for drilling hazards and failure rates is the lease–unit where the initial time is the date of leasing and the failure time is the date a well is drilled in that unit. Since there are multiple leases per unit, I downweight each lease–unit overlap by its area so that the shares of mineral ownership in a unit sum to one.

[4] This is even more evident in the cumulative failure rate, shown in Figure 19 in the Appendix.

Table 1: Well production and length

	Log monthly production			Log well length
	OLS	Section FE	Well FE	OLS
$\log \min\{t, t_{int}\}$	-0.53^{***}	-0.53^{***}	-0.53^{***}	
	(0.02)	(0.02)	(0.02)	
$\max\{\log t - \log t_{int}, 0\}$	-1.32^{***}	-1.33^{***}	-1.34^{***}	
	(0.01)	(0.01)	(0.01)	
Log well length	0.99^{***}	0.75^{***}		
	(0.12)	(0.13)		
Log OGIP	0.52^{***}			0.01
	(0.05)			(0.06)
Is 1st well drilled in section	-0.07	0.06		-0.02
	(0.04)	(0.06)		(0.02)
Spud date (years since 2000)	0.05^{***}	0.01		-0.03
	(0.01)	(0.02)		(0.02)
Is cross-unit well				0.35^{***}
				(0.06)
Blended royalty rate	3.73^{***}			0.22
	(0.51)			(0.20)
Fixed effects	*No*	*Section*	*Well*	*No*
Num. sections (i)	1109	1109	1109	1109
Num. section-wells (iw)	1874	1874	1874	1874
Num. section-well-months (iwt)	$100,982$	$100,982$	$100,982$	

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. Standard errors clustered by section. Intercepts omitted.

Three simple regressions and plots of drilling hazards provide clues about how unobservable geological productivity may affect royalty rates, drilling decisions, and production outcomes. Table 1 shows regression estimates of the logarithm of natural gas production from well w in section i in month t using several covariates. Model (1) is simple OLS. Model (2) includes section-specific (i) fixed effects that control for unobservable productivity shared by the wells in each section. Model (3) uses finer-grained well-specific (iw) fixed effects. All three models cluster standard errors at the section level, thereby accounting for serial correlation within wells iw and correlation between wells within a section i.

OLS estimates in Model (1) suggest that initial wells (those for which $w = 1$) produce less on average than later wells, even if we condition on observable factors. We can see this by inspecting the coefficient on the indicator variable "Is 1st well drilled in section." There are two explanations for the negative coefficient. First, firms might only drill multiple wells in locations they find to be especially productive. This is simply upward selection bias. Second, firms may be drilling initial low-cost, low-productivity wells simply to hold expiring leases. Model (2) uses fixed effects at the section (i) level to eliminate selection on section-specific productivity and focus only on this second, low-cost initial well explanation. The estimate on the "Is 1st well drilled in section" coefficient is no longer statistically significant in Model (2) and, in fact, flips sign. Thus, there is no statistical support for the hypothesis that Well 1 is different from Wells 2+ either because of lower effort or worse technology. Ruling such possibilities is important if we are to focus solely on firms' decision of *whether* to drill and ignore to their decisions of *how* to drill.

Model (1) also shows us that wells on sections with higher average royalty rates produce more. To call this relationship causal would be absurd, particularly since high royalty rates should reduce the returns to a firm's drilling effort and, hence, production. There are two possible explanations for this positive correlation. First, when royalty rates are high, only more productive locations will be profitable and, hence, drilled. This would imply upward selection bias on the royalty-rate coefficient. Second, royalty rates may be be positively correlated with a firm's prior beliefs about geological productivity. We can also see evidence for the royalty–productivity correlation in Figure 5. The figure plots an estimate of the probability that a first well is drilled within a given number of months (the failure function), and it breaks this out by royalty rate. With the notable exception of the few locations that have 12.5% royalty rates, the probability of being drilled by a particular date increases with the royalty rate. This means that firms are accelerating investment in locations with higher royalty rates compared to locations with lower royalty rates. Thus, not only are these high royalty rate locations more productive, they are also more *profitable*—a reversal which is much harder to generate with only the first selection bias explanation. Royalty rates are therefore likely to be positively correlated with unobserved productivity through firms' initial beliefs about geological quality. Moreover, these beliefs are likely to be informative. When firms believe an area to be more productive, they pay higher royalty rates and accelerate investment, and the wells produce more.

The observable geology component, oil and gas in place (OGIP), has qualitatively similar impacts to royalty rates in Model (1). The highly significant, positive coefficient on OGIP implies that wells located on observably better geology produce more, just like wells with royalty rates.

A serious concern associated with modeling firms' decisions to drill as simple, discrete choices is that firms might be exerting different amounts of effort based on the profitability of each location. The above discussion hints at this in relation to whether initial wells are low-effort and, hence, see low production, but the concern is relevant for other observable factors that influence productivity. Fortunately, such concerns do not appear to be warranted in my sample. To further investigate whether firms drill initial wells differently, we can turn to the most obvious measure effort: the length of each well's horizontal component. As long as marginal well production diminishes with well length, producers should exert more effort and drill longer wells in higher-productivity locations. To test whether operators drill longer wells in better locations, Model (4) regresses the logarithm of well length on observable characteristics of the sections and wells. Observable

Figure 5: Failure function for initial drilling by geology, royalty

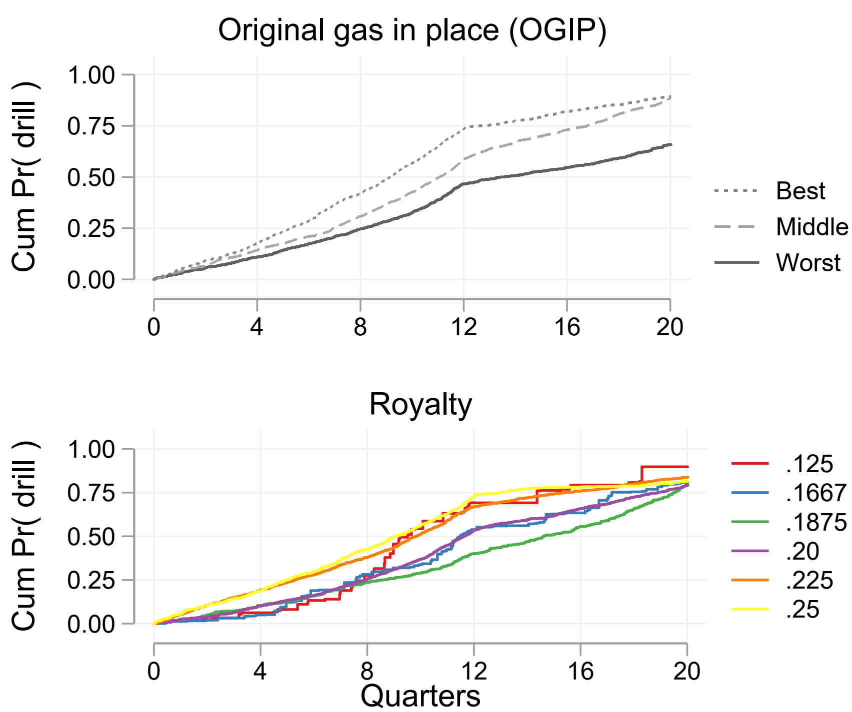

geological quality (OGIP), unobservable quality as captured by royalty rates, the date the well is drilled, and the order of the well in a section all have no statistically discernible effect on well length. The only significant effect on well length is whether a well has a permit to produce hydrocarbons in multiple drilling units, something only a few of the wells in my sample have. Instead, the length of most wells instead appears to be limited by the size of each section. The violin plots in Figure 6 break out the distribution of well-lengths for wells that lack cross-unit permits by the order the well was drilled in a section. There is clearly a mode just shy of a mile long (5,280 feet)—the width of a section—and the tight distributions suggest that regulatory constraints drive well design decisions.

Figure 6: Distribution of well-length

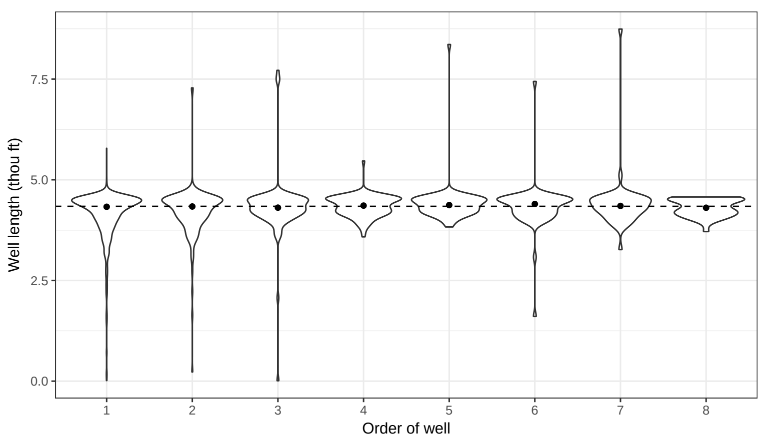

Dots mark category specific medians. Dashed lines mark sample median. Sample excludes cross−unit wells.

A final concern that must be addressed is whether the production technology improves over the course of my sample. As discussed in the introduction, a good reason to study the Haynesville during the period I do is that technology and process improvements may have had less of an impact in the Haynesville compared to other shale plays. To address this, I include each well's spud date (the date drilling started) in years since 2000 in Models (1) and (2) of well production. For example, for a well drilled in July 2007, this variable would take the value 7.5. The coefficient is called "Spud date (years since 2000)." In Model (1), the coefficient is positive and significant, suggesting a 5% annual improvement in the productivity of wells drilled. Over the course of eight years, this estimate would suggest an improvement in productivity of around 39%. Nevertheless, the estimate is suffers from upward selection bias, as sections that continue to see drilling in later years are more likely to be the most productive ones. Model (2) eliminates any such selection bias by including section-specific fixed effects and relying on within-section variation of well timing to identify upward productivity trends. While the coefficient remains positive, we cannot reject that it is zero at any conventional level of significance.

6 Model

To understand whether learning about geology affects how average output per well evolves, we need to know how firms' beliefs about the geological productivity in each location of a section evolve over time. I assume a simple information structure and use an econometric model to recover the joint joint distribution of beliefs before and after drilling.

The sequence of events in the model, depicted in Figure 7, is as follows. Upon arriving at a section to negotiate a lease, firms receive an initial, noisy signal about the section's quality. This is their prior. While we cannot observe the firm's signal, the signal can affect the outcome of negotiations over royalty-rates, as well as the firm's eagerness to start drilling. Once the firm drills an initial well, it learns the true quality of the section. Knowing this, the firm then decides if and when to drill up to eight more wells. The section's true quality determines the volume of production for each well.

Sections are indexed by $i = 1, \ldots, N$, and wells on a section are numbered by the order in which they are drilled $w = 1, 2, \ldots, 8$. Time is indexed by t and measured in months. I index the leases in a given section by j. Where possible, I use lower-case letters to denote specific realizations of random variables and upper-case letters to denote the random variables themselves. One exception to this is the variable D_{it}, which is the cumulative number of wells drilled prior to time t.

Table 2: Ordered probit regression: royalty rates

	Model 1
Log OGIP	0.08
	(0.08)
Share of out-of-state grantors	1.13^{***}
	(0.16)
Log median house value	0.56^{***}
	(0.08)
Share of permeable surface	-1.19^{*}
	(0.55)
0.125|0.1667	3.65^{***}
	(1.10)
0.1667|0.1875	4.01^{***}
	(1.10)
0.1875|0.2	4.78^{***}
	(1.10)
0.2|0.225	5.69^{***}
	(1.10)
0.225|0.25	6.33^{***}
	(1.10)
AIC	4247.20
BIC	4294.49
Log Likelihood	−2114.60
Deviance	4229.20
Num. obs.	1414

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$

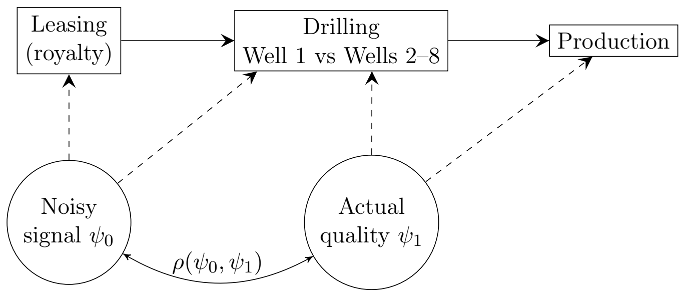

Figure 7: Correlated initial signal (ψ_0) and actual quality (ψ_1) link outcomes

6.1 Information

The geological productivity of a section has two components. The first is based on public information, which I assume to be the OGIP measure in Browning et al. (2015) and Gülen et al. (2015). The second is an unobserved component orthogonal to OGIP that the firm, but not the econometrician, knows. For section i, denote the realization of this second component as ψ_{i1}. A firm's prior belief about it is ψ_{i0}. The prior belief, ψ_{i0}, affects initial royalty-rate negotiations and initial drilling. Actual unobservable quality, ψ_{i1}, affects firms' subsequent drilling decisions and the realized production outcomes from each well. Figure 7 illustrates this information structure. The dashed lines show how initial signal (ψ_0) and actual quality (ψ_1) impact three observable outcomes shown in boxes: royalty rates, drilling decisions, and production. I assume that signal and true quality are jointly normal variables.[5]

Assumption 1 *The initial signal and true quality are jointly distributed as a bivariate standard normal variable with correlation ρ:*

$$\begin{pmatrix} \psi_{i0} \\ \psi_{i1} \end{pmatrix} \sim N\left(0, \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix}\right).$$

[5] Setting the variance to 1 is simply a normalization as ψ_0 and ψ_1 will enter each equation with scalar coefficients. Moreover, it is the *correlation* of ψ_0 and ψ_1, not the relative size of their variances that matter. One could, in fact, simply interpret ψ_0 as ψ_1 plus an independent, normally distributed noise term, scaled by the variance of ψ_1 plus noise. This is, in fact, how ψ_0 is constructed in estimation.

The correlation between them, ρ, is a key parameter. It measures how informative firms' initial beliefs are and thereby determines the capacity for learning about geology to impact average output per well. Given Assumption 1, the conditional distribution of ψ_{i1} given ψ_{i0} is simply

$$F(\psi_1|\psi_0) = N\Big(\rho\psi_0,\ (1-\rho^2)\Big).$$

6.2 Royalty rates

Royalty rates are the outcome of a one-time negotiation between mineral owners and firms. Since we know little about the information structure of the game being played between owners and firms, I model the outcome in a way that allows (but does not require) firms' information to affect the royalty rate. This means I do not have to make potentially invalid assumptions about what mineral owners know about the actual quality of their minerals.

A royalty rate in section i is a discrete random variable $R_i \in \{\bar{r}_1, \ldots, \bar{r}_L\}$. It is determined by a continuous latent variable R_i^*, which is a linear combination of economic variables that affect firms' willingness to pay for mineral rights and ones that affect mineral owners' willingness to accept a firm's offer. Two variables affect a firm's willingness to pay: public information about section i's geology, G_i, and the firm's initial signal, ψ_{i0}. Mineral owner characteristics, X_{ri}, affect owners' bargaining position but cannot directly affect firms' profits except through the royalty rate. This is an an identifying exclusion restriction: it rules out the possibility that landowners with low willingness to accept drilling impose restrictions that affect firms' drilling costs. Finally, the latent R_i^* is affected by a random bargaining shock, ν_i. We can write a realization of R_i^* as

$$r_i^* = \underbrace{\beta_\psi \psi_{i0} + \beta_g g_i}_{\text{Firm WTP}} + \underbrace{\beta_x^\top x_{ri}}_{\text{Landonwer WTA}} + \underbrace{\nu_i}_{\text{Barg. shock}}. \tag{1}$$

Royalty rates take a discrete value r_l when r_i^* falls between two corresponding thresholds:

$$r_i = \bar{r}_l \iff \kappa_{l-1} < r_i^* \leq \kappa_l \tag{2}$$

where $-\infty = \kappa_0 < \ldots < \kappa_l < \ldots < \kappa_L = +\infty$.

To form the likelihood of observing a particular royalty rate, I assume that ν_i and ψ_{i0} are independent normal variables. This implies royalty rates can be modeled with an ordered probit regression that includes a random effect.[6] Setting the variance of ν_i to one is just a normalization.

Assumption 2 *ψ_{i0} and ν_i are independent, normal random variables. That is, $f(\nu_i|g_i, x_{ri}, \psi_{i0}) = \phi(\nu_i)$, where $\phi(\cdot)$ is the PDF of the standard normal distribution.*

Given Assumption 2 and denoting the CDF of the standard normal distribution as $\Phi(\cdot)$, the likelihood of observing a particular royalty rate $r_i = \bar{r}_l$ given x_i can be written as

$$\begin{aligned} L_i(R_i = \bar{r}_l|\psi_{i0}, g_i, x_{ri}) = \Phi&\left(\kappa_l - \beta_\psi\psi_{i0} - \beta_g g_i - \beta_x^\top x_{ri}\right) \\ &- \Phi\left(\kappa_{l-1} - \beta_\psi\psi_{i0} - \beta_g g_i - \beta_x^\top x_{ri}\right). \end{aligned} \quad (3)$$

6.3 Drilling decision

The central part of my statistical model is a firm's monthly choice to drill a discrete number of wells. I model this decision using a Rust (1987)dynamic discrete choice framework. In each month t and each section i, the operator decides how many wells to drill: $d_{it} \in \{0, 1, 2, \ldots, \bar{d}\}$. Based on my data, I set $\bar{d} = 8$ wells per quarter. The operator's decision affects the firm's ability to drill in the future and, if the firm has not drilled before, its information.

The state that determines the set of firms' choices is s_{it}. This endogenous state variable includes information about the months remaining until a lease's primary term expires, the months remaining until its extension also expires, and the cumulative number of wells drilled to date. I denote this last component, cumulative prior drilling, as D_{it}. The firm is not able to drill once the primary term expires. Similarly, its total drilling is limited to

[6] While it would theoretically be useful to model lease-specific royalty rates with a unit-level random effect, the vast majority of variation in royalty-rates is at the unit level. This means that the unit-specific random effect swamps lease-specific variation and makes the model numerically unstable.

a maximum total of eight wells. I denote the firm's action space as correspondence Γ:

$$\Gamma(s_{it}) = \begin{cases} \{0\} & \text{if lease extension expired or } D_{it} = 8 \\ \left\{0, 1, \ldots, \min\left\{\bar{d}, 8 - D_{it}\right\}\right\} & \text{otherwise} \end{cases}.$$

The firm receives at least one and possibly two realizations of its information about unobserved geological quality, ψ. The first, ψ_{i0}, is noisy, and all firms receive it before they make any drilling choices. Firms elect whether to learn the true quality, ψ_{i1}, by choosing drill an initial well. The transition of the firm's information can be expressed as

$$\psi_{i,t+1} \sim \begin{cases} N\left(\rho\psi_{it},\ (1-\rho^2)\right) & \text{if } D_{it} = 0 \text{ and } d_{it} > 0 \\ \psi_{it} & \text{otherwise} \end{cases}.$$

When making decisions, firms also take into account a vector of observable state variables, Z_{it}. These have exogenous transitions and affect firms' payoffs to drilling. We can group them into two components. The first component, Z_{1it}, contains market price signals (natural gas prices, costs, and a year effect). I assume that these follow an first-order Markov process. The fact that these prices are set in a large, national market justifies the assumption of exogeneity. The second set of state variables, Z_{2i}, is time-invariant and contains the average royalty-rate, R_i and the observable component of geology, G_i. While Assumption 3 below implies that the transition of Z_{it} is independent of ψ_{it}, it does not rule out dependence between Z_{it} and ψ_{it} because R_i may depend on ψ_{i0}.

Assumption 3 *The vector $Z_{i,t+1}$ is conditionally independent of the other state variables: $F_z(Z_{i,t+1}|z_{it}, s_{it}, \psi_{it}, \epsilon_{it}, d_{it}) = F_z(Z_{i,t+1}|z_{it})$*

Finally, each period, the firm also receives a random, $\bar{d}$-length vector of profitability shocks, ϵ. Each component of ϵ is associated with a particular choice of the number of wells to drill, d. These shocks could include, for example, weather disruptions or availability of a suitable rig in the local area.

6.4 Flow payoffs

Firms' decisions to drill are based on the sum of static and dynamic payoffs from drilling. I assume that static profits are simply revenues less costs. Given a choice to drill d wells in section i in period t, the net (flow) payoff to drilling can be written as

$$u_d(z_{it}, s_{it}, \psi_{it}, \epsilon_{it}) = \mathbb{E}\left[revenue(d, z_{it}, s_{it}, \psi_{i1}) | z_{it}, s_{it}, \psi_{it}\right] - cost(d, z_{it}, s_{it}, \epsilon_{it}). \tag{4}$$

The fact that unobserved heterogeneity affects revenues, not costs, reflects the idea that heterogeneity is in geological prodictivity.[7] Revenues are essentially the number of wells drilled, d times the value of an additional unit of total production, p_t, times the Expected Ultimate Recovery (EUR) of each well:

$$revenue(d, z_{it}, s_{it}, \psi_{i1}) = d(1 - r_i) p_t Q(g_i, \psi_{i1}). \tag{5}$$

EUR is calculated differently depending on whether the firm has drilled ($D_{it} > 0$) and knows ψ_{i1} or whether the firm has not ($D_{it} = 0$) and must take a conditional expectation given the noisy signal, ψ_{i0}:[8]

$$Q(g_i, \psi_{i1}) = \exp\left\{\alpha_0 + \alpha_g g_i + \alpha_\psi \psi_{i1}\right\} \tag{6}$$

$$\mathbb{E}[Q(g_i, \psi_{i1}) | \psi_{i0}, g_i] = \exp\left\{\alpha_0 + \alpha_g g_i + \alpha_\psi \left(\rho \psi_{i0}\right) + \alpha_\psi^2 \frac{(1 - \rho^2)}{2}\right\} \tag{7}$$

When evaluating the financial profitability of a well, what firms care about is not the current price of natural gas, but the present value of the price at which the gas will be sold when it is produced. Operators often sell gas production forward, hedging against future price drops and locking in

[7] This correlation is also reflected in a long history of papers which analyze oil and gas auctions in a common-values paradigm.

[8] The joint normality of ψ_{i1}, ψ_{i0} and their independence from g_i and p_t imply the form of the conditional expectation.

revenues when production commences.[9] Thus, I use a weighted average of the forward curve that incorporates both well decline and time-discounting to capture firms expected production revenue. Let $F(t, t+\tau)$ be the monthly average futures price at time t for gas delivered at time $t+\tau$ where both t and τ are measured in months. Following Covert (2015), I assume that a shale gas well produces for 20 years. Then the relevant gas price for the firm is a weighted and discounted average of futures prices:

$$p_t = \frac{\sum_{\tau=1}^{240} \tilde{\beta}^{\frac{\tau}{12}} \exp\{f(\tau;\hat{\gamma}, t_{int})\} F(t, t+\tau)}{\sum_{\tau=1}^{240} \exp\{f(\tau;\hat{\gamma}, t_{int})\}} \tag{8}$$

where $\tilde{\beta}$ is the nominal discount factor, $f(\tau;\hat{\gamma}, t_{int})$ is expected production decline curve with parameter estimates $\hat{\gamma}$ taken from fixed effects estimates in Table 1. The variable p_t then represents the marginal value of an additional unit of expected ultimate recovery (EUR).

Reliable measures of forward prices, $F(t, t+\tau)$, are only available for τ up to 5 years. To account for this, I replace $F(t, t+\tau)$ for years 6–24 with the average 5-year futures price, $\overline{F(t, 5\ year)} = \frac{1}{12}\sum_{m=1}^{12} F(t, 48+m)$. Rather than attempt to estimate β, I set it exogenously as is typical in empirical dynamic discrete choice papers. I follow Kellogg (2014), who assumes a nominal discount rate of 12.5% based on a survey of the Society of Petroleum Evaluation Engineers. I also compute average inflation from the average change in the logarithm of the PPI for final goods less energy and food over the sample period Jan 2003–Oct 2016. This is 2.34%. Combining the two, this gives me an annual nominal discount factor of $\tilde{\beta} = 1/1.125 \approx 0.89$ and an annual real discount factor of $\beta = 1.0234/1.125 \approx 0.91$, which is close to the value 0.9 used by Covert (2015) and Muehlenbachs (2015) for similar applications, as well as the real discount rate used in Kellogg (2014).

The drilling cost function is a function of the number of wells (d), whether the firm has to sign a lease extenion and pay the mineral owner again ($\mathbb{1}$[sign ext]), the drilling cost index (c_t), the date (t), and a $\bar{d}$-length vector of i.i.d. cost shocks, ϵ. Rather than include a large number of time fixed ef-

[9] One could also justify this by assuming that the futures market accurately reflects firms' expectations about future prices.

fects, I include a third-order Chebyshev polynomial of the year. This allows some flexibility for the model to match adjustment costs at the beginning of the boom while still being parsimonious. To make the model stationary while still allowing firms to anticipate future costs, I assume that there is a 1/4 probability that the year will increase next quarter. I then assume that trend does not change after 2016. I also allow drilling costs per well to differ depending on whether the firm drills just one ($\alpha_{c_{cheb0,1}}$) or multiple ($\alpha_{c_{cheb0,2+}}$).

$$cost(d, s_{it}, \epsilon_{it}) = \begin{cases} \alpha_{exttension} \mathbb{1}[\text{sign ext}] - \epsilon_{it}(0) & \text{if } d = 0 \\ 1 \left[\alpha_{c_{cheb0,1}} \quad + p(t; \alpha_{cheb}) + \alpha_{drilling} c_t\right] - \epsilon_{it}(1) & \text{if } d = 1 \\ d \left[\alpha_{c_{cheb0,2+}} + p(t; \alpha_{cheb}) + \alpha_{c_{drilling}} c_t\right] - \epsilon_{it}(d) & \text{if } d > 1 \end{cases} . \tag{9}$$

6.5 The Value Function

Given a discount rate is $\beta \in (0, 1)$, a firm's objective is to choose a sequence of actions, $\{d_{it}\}_{t=1}^{\infty}$, to maximize the sum of its expected future payoffs, $\tilde{u}$:

$$V(s_{i0}, z_{i0}, \psi_{i0}, \epsilon_{i0}) = \max_{\{d_{it}\}_{t=0}^{\infty}} \mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t \tilde{u}(d_{it}, S_{it}, Z_{it}, \psi_{i1}, \epsilon_{it}) \middle| s_{i0}, z_{i0}, \psi_{i0}, \epsilon_{i0}\right]$$

$$\text{subject to}$$

$$d_{it} \in \Gamma(S_{it}).$$

The action-space, $\Gamma(s) \subseteq \{0, 1, \ldots, \bar{d}\}$, is compact, so under the assumption that $\tilde{u}$ is bounded and continuous, the maximum exists. We can write the firm's problem as a functional equation:

$$\begin{aligned} V(s_{it}, z_{it}, \psi_{it}, \epsilon_{it}) = & \max_{d \in \Gamma(s_{it})} u_d(s_{it}, z_{it}, \psi_{it}, \epsilon_{it}) \\ & + \beta \, \mathbb{E}\left[V(S'(s_{it}, d), Z_{i,t+1}, \Psi_{i,t+1}, \epsilon_{i,t+1}) \middle| s_{it}, z_{it}, \psi_{it}, \epsilon_{it}, d\right]. \end{aligned}$$

There are two absorbing states: when a lease but the firm hasn't drilled, and when the firm exhausts all drilling opportunities. In both cases, the firm can no longer drill new wells, so I assume that the value of being in these states is zero. This provides the following terminal condition:

$$V(s_{it}, z_{it}, \psi_{it}, \epsilon_{it}) = 0 \text{ if lease expired or } D_{it} = \bar{D}.$$

Solving the value function can be done using backwards recursion starting from the point at which all eight wells have been drilled.

The cost function, equation 9, imposes the assumption below that choice-specific shocks, ϵ_d, are additively separable from static payoffs. I also assume below that they are conditionally independent from the other state variables. These are both standard assumptions for dynamic discrete choice models.

Assumption 4 *Flow payoffs are additively separable with respect to the dth dimension of the choice-specific shock:*

$$u_d(s, z, \psi, \epsilon) = u_d(s, z, \psi) + \epsilon_d$$

Assumption 5 *The joint density of the state variables can be factored as*

$$\begin{aligned} &f(s_{i,t+1}, z_{i,t+1}, \psi_{i,t+1}, \epsilon_{i,t+1} | s_{it}, z_{it}, \psi_{it}, \epsilon_{it}) = \\ &\qquad f_\epsilon(\epsilon_{t+1} | s_{t+1}, z_{t+1}, \psi_{i,t+1}) f(s_{t+1}, z_{i,t+1}, \psi_{i,t+1} | s_{it}, z_{it}, \psi_{it}, d_{it}) \end{aligned}$$

As previously mentioned, Assumption 5 does *not* imply that unobserved heterogeneity in geological quality, ψ_{it}, is independent of z_{it}. Royalty rates are allowed to depend on ψ_{i0}, but only through equations (1) and (2).

Continuing to follow the literature on dynamic discrete choice models, I work with the integrated value function, (also called an expected value function or $\mathbb{E}\max$ function). Dropping the i subscript and denoting $t+1$ with a trailing $'$ and $t+2$ with a trailing $''$, we can write the integrated value function as

$$\mathbb{E}\mathrm{V}(s', z, \psi) = \mathbb{E}\left[\max_{d \in (s')} \left\{ u_d(s', Z', \Psi') + \epsilon'_d + \beta\, \mathbb{E}\left[V(S'', Z'', \Psi'') \middle| d, s', Z', \Psi' \right] \right\} \middle| z, \psi \right].$$

I then define the choice-specific (alternative-specific) value function v_d as

$$v_d(s,z,\psi) = u_d(s,z,\psi) + \beta\, \mathbb{E}\mathrm{V}(S'(s,d),z,\psi).$$

The choice-specific value function plus shock $v_d + \epsilon_d$ are the expected economic payoff that the firm receives upon drilling. Thus, we can take an expectation over ϵ and write the integrated value function as

$$\mathbb{E}\mathrm{V}(s,z,\psi) = \mathbb{E}\left[\max_{d\in\Gamma(s)} \{v_d(s,z,\psi) + \epsilon_d\}\right].$$

Given a vector of parameters that characterize payoffs u_d, it is straightforward to compute the integrated value function using a combination of value-function and policy-function iteration.

To form the likelihood, I assume that vector of choice-specific shocks ϵ is composed of random draws from a multivariate Type-I Extreme Value (Gumbel) distribution. Following my convention of differentiating between random variables and their realizations where possible, I denote specific realizations of ϵ as ε.

Assumption 6 $\epsilon_d \sim_{iid}$ *Gumbel*$(-ec, 1)$. *That is,*

$$f_\epsilon(\varepsilon_1,\ldots,\varepsilon_{\bar{d}}) = \prod_{d=1}^{\bar{d}} \exp\left\{-\exp\left\{-[\varepsilon_d - ec]\right\}\right\}$$

where ec is the Euler–Mascheroni constant.

Assumptions 5 and 6 imply that choice-specific shocks represent idiosyncratic events like a rig or extra materials becoming available nearby. Serial correlation is explicitly ruled out. This implies there are no permanent shifts in the firm's signal about productivity from new information except for the update from drilling an initial well. Instead, serial correlation in profitability is captured exclusively through ψ_{it}.

Given the above assumptions, the integrated value function has the standard closed-form:

$$\mathbb{E}\mathrm{V}(s,z,\psi)=\log\sum_{d\in\Gamma(s)}\exp\left\{v_d(s,z,\psi)\right\},$$

and the probability of observing action d conditional on all state variables except action-specific shocks ϵ is

$$\Pr\left(d|s,z,\psi\right)=\frac{\exp\left\{v_d(s,z,\psi)\right\}}{\sum_{l\in\Gamma(s)}\exp\left\{v_l(s,z,\psi)\right\}}.$$

The fact that there are multiple leases and, therefore, multiple expiration dates per section complicates estimation. I assume that only one of the lease-expirations matters to the firm. The probability that the firm picks a particular lease j and its expiration is proprtional to Y_{ij}, the logarithm of lease size. The deterministic transition of s_{it} further implies that Y_{ij} only affects decisions through the initial s_{ij0}, so that

$$\Pr(s_{ij,t+1}|s_{ijt},y_{ij},d_{it})=\Pr(s_{ij,t+1}|s_{it},d_{it}).$$

The likelihood of observing a sequence of decisions d_{it} for an entire unit given the noisy signal ψ_{i0} and true ψ_{i1} (as well other variables) can then be written as

$$\begin{aligned}L_i(\{d_{i,t+1},s_{i,t+1}\}|\{z_{it}\}_{t=1}^{\bar{T}_i},\{y_{ij}\}_{j=1}^{J_i},\psi_{i0},\psi_{i1})=&\\ \left[\prod_{t=T_{1i}+1}^{\bar{T}_i}\Pr(d_{i,t+1}|s_{i,t+1},z_{i,t+1},\psi_{i1})\Pr(s_{i,t+1}|s_{it},d_{it})\right]&\\ \times\left[\sum_{j=1}^{J_i}\left(\prod_{t=1}^{T_{1i}}\Pr(d_{i,t+1}|s_{ij,t+1},z_{i,t+1},\psi_{i0})\Pr(s_{ij,t+1}|s_{ijt},d_{it},j)\right)\Pr\left(j|y_{ij}\right)\right].&\end{aligned}\tag{10}$$

6.6 Production

The final component of the model consists of monthly production outcomes from each well. These outcomes function as a set of measurements of true unobserved productivity, ψ_{i1}. I assume the following econometric model of monthly production for well w in section i at time t:

$$\log q_{iwt} = \gamma_0 + \underbrace{f(t;\gamma,t_{int})}_{\text{decline}} + \underbrace{\alpha_g g_i}_{\text{unit-specific}} + \underbrace{\gamma_{len}\log length_{iw}}_{\text{well-specific}} + \xi_{iwt}$$

$$\xi_{iwt} = \alpha_\psi \psi_{i1} + u_{iw} + \eta_{iwt} \tag{11}$$

where

$$f(t;\gamma,12) = -\gamma_1 \log\min\{t,12\} - \gamma_2(\max\{\log t - \log 12, 0\}) \tag{12}$$

and the unobserved components are independent normal variables:

$$u_{iw} \sim_{iid} N(0,\sigma_u^2) \qquad \eta_{iwt} \sim_{iid} N(0,\sigma_\eta^2).$$

The term $f(t;\gamma,12)$ captures exogenous, natural production decline over time. The function f is allowed to have a kink at 12, though it will still be continuous. This is equivalent to estimating a traditional Arps model of decline in which a well's flow-rate is $dQ/dt = Q(0)t^\gamma$ and there is a break in γ at $t = 12$ months.[10] There are two sets of observed, time-invariant variables that determine production. The first, g_i, captures the effect of observable geology and is shared by all wells within the unit. The second, $\log length_{iw}$ is

[10] Patzek, Male, and Marder (2013) analyze the physics behind shale well decline curves and shows that two physical processes determine decline rates. During the first phase, which occurs during the interval $t \in (0, t_{int})$, cumulative production $Q(t)$ can be modeled as $Q(t) = \mathcal{K}\sqrt{t}$. This implies that $dQ/dt = 0.5\mathcal{K}t^{-0.5}$. After the interference time (t_{int}), the authors find that production should follow an exponential decline where $dQ/dt = q_0 e^{-\delta t}$. Male et al. (2015) find that the interference time t_{int} for Haynesville shale wells is around a year. While I confirm that a 12-month breakpoint appears to fit the data best compared to months 3–59, I find that the second regime is best fit using a linear decline rate: when including both a linear time trend and $\log t$ in a model of the logarithm of monthly production, the linear time trend is much less significant than $\log t$ during the second regime. Thus, I drop the linear term for regime two and keep only $\log t$ in regime two.

well-specific and captures the length of the well. The unobserved component, ξ_{iwt}, is decomposed into three parts: the unobserved quality of the section, ψ_{i1}, a well-specific term the firm cannot forecast, u_{iw}, and a random shock to monthly production, η_{iwt}. The coefficients α_g and α_ψ are the same as those from equations 6 and 7 in the revenue equation for firms' drilling payoffs. This restriction is consistent with rational behavior by firms, and it also helps identify firms' payoffs.

I assume that that u_{iw} and η_{iwt} are normally distributed, uncorrelated random effects. This is made explicit in Assumption 7:

Assumption 7 *The unobserved components u_{iw} and η_{iwt} are i.i.d. normally-distributed variables*

$$f(u_{iw}, \eta_{iwt} | \psi_{i1}, g_i, \log length_{iw}) = \phi\left(\frac{u_{iw}}{\sigma_u}\right)\phi\left(\frac{\eta_{iwt}}{\sigma_\eta}\right).$$

When there are T_{iw} months of production data for well w in section i, we can use Assumption 7 to write the likelihood of $\log q_{iwt}$ conditional on ψ_{i1} and other observables as

$$\begin{aligned} L\left(\log \vec{q}_{iw} | \psi_{i1}, g_i, \log length_{iw}\right) = \\ -\frac{1}{2}\left[T_{iw}\log(2\pi) + (T_{iw}-1)\log\sigma_\eta^2 + \log(\sigma_\eta^2 + \sigma_u^2 T_{iw})\right] \\ -\frac{1}{2\sigma_\eta^2}\left[\sum_{t=1}^{T_{iw}}(u_{iw}+\eta_{iwt})^2 - \frac{\sigma_u^2}{\sigma_\eta^2+\sigma_u^2 T_{iw}}\left(\sum_{t=1}^{T_{iw}}(u_{iw}+\eta_{iwt})\right)^2\right] \end{aligned} \tag{13}$$

where the $u_{iw} + \eta_{iwt}$ is defined according to equation (11).

6.7 Model likelihood

We can write the likelihood conditional on the signal, ψ_{i0}, and true quality, ψ_{i1}, as the product of the likelihood of the royalty rate, the history of drilling decisions, and production from the wells that may have been drilled:

$$L(history_i|\psi_{i0}, \psi_{i1}) = L(r_l|\psi_{i0})\ L(\{d_{it}\}|\psi_{i0}, \psi_{i1})\ \prod_{d=1}^{D_i} L\left(\{\log q_{iDt}\}_{t=2}^{T_{qiD}}\middle|\psi_{i1}\right). \tag{14}$$

Since ψ_{i0} and ψ_{i1} are not observed, I integrate them out by simulation.[11] Given M draws of ψ_{i0}, ψ_{i1}, the simulated likelihood is then

$$SL(history_i) = \frac{1}{M}\sum_{m=1}^{M} L_i\left(history_i|\psi_{im0}, \psi_{im1}\right). \tag{15}$$

To form the complete likelihood of the data I observe, I assume conditional independence of section histories given exogenous prices, observed geology, and lease characteristics.

Assumption 8 *Choice specific shocks (ϵ_{it}), royalty-rate shocks (ν_i), and lease-shocks (ζ_{ji}) are independent across sections.*

This rules out within-firm profitability shocks such as a firm waiting to drill until a rig on a neighboring section is free and firms drilling cross-unit wells. It also rules out leases having interest in more than one section, which is generally the case. Additionally, I make the following assumption, which rules out spatial dependence:

Assumption 9 *Signals and quality are uncorrelated across sections i:*

$$f\left((\psi_{10}, \psi_{11}), \ldots, (\psi_{I0}, \psi_{I1})\right) = \prod_{i=1}^{I} f(\psi_{i0}, \psi_{i1}).$$

Assumption 9 rules out spatial correlation between unobserved quality in sections. This implies that there are no information externalities wherein wells in one section are informative about wells in a neighboring section. This assumptions about spatial dependence is not likely to fully hold in

[11] Given a correlation ρ, I draw two independent standard normal variables v_{1i} and v_{2i} and form $\psi_{i1} = v_{1i}$ and $\psi_{i0} = \rho v_{1i} + \sqrt{1-\rho^2}v_{2i}$. This is simply multiplying a bivariate, independent normal variable by the Cholesky decomposition of its covariance matrix.

practice. However, allowing for spatial correlation greatly complicates the model as firms' optimal decisions must be made over sets of drilling units and the order in which they are drilled.[12] The full, simulated likelihood of the data is therefore

$$SL(data) = \prod_i SL(history_i), \tag{16}$$

and I numerically maximize it to obtain estimates.

6.8 Identification

Given the above assumptions, the model is statistically identified. Intuitively, firms that pay high royalty rates and accelerate initial drilling are likely to have received high initial signals. Sections with high actual quality should see firms start intense, additional drilling after the initial well, and the wells drilled should produce large volumes of gas.

The royalty-rate equation serves as a single measurement equation for the firm's signal during the initial drilling phase. I assume that the characteristics of surface land and mineral owners (x_{ri} in equation (1)) only affect firms' profits through mineral owners' bargaining position in royalty rate negotiations. Timmins and Vissing document that higher socio-economic status households have more leverage in negotiations with landmen (Timmins and Vissing 2014; Vissing 2015, 2016), and Hitaj, Weber, and Erickson (2018) documents that absentee mineral owners behave differently than local mineral owners in leasing rural acreage. Based on these findings, I include median housing values, the imperviousness of a location's surface (a measure of urbanization), and the share of minerals owned by out-of-state individuals as exogenous bargaining shifters in x_{ri}. Note that time-varying variables do not enter this equation because it is the *blended* royalty rate—an average over all leases in a section—that matters. Thus the point of time associated with a royalty rate is not well-defined.

[12] With a larger set of data, one solution would be to draw a sample of drilling units such that no drilling unit touches another.

Firms' drilling decisions are affected by exogenous variation in price levels and volatility. As mentioned, these are determined in a national market, and a single well's production will not move the market. Changes in prices provide variation over time and, along with the very strong structure on the deterministic state-transitions in the drilling decision problem, helps explain time variation in decisions.

Finally, the production equation provides measurement equations for a location's true quality, and pins down the contribution of observed and unobserved geology to firm's revenue: α_g and α_ψ. We observe variation in well-length, observable geology g_i, and the amount of time a well has produced. Many sections see multiple wells being drilled. Under Assumption 7 that production follows an additive random effects model, the distribution of ψ_{i1} is immediately identified, and estimation is possible through generalized least-squares or maximum likelihood. This pins down the distribution of $\alpha_\psi \psi_{i1} + u_{iw}$. The presence of multiple wells then identifies γ_ψ and σ_u.

7 Results

I estimate the parameters of the value function using the standard Rust (1987) Nested Fixed Point (NFXP) algorithm in which I maximize a simulated likelihood (MSL).[13] I first discuss the specifics of the royalty-rate and production models and then proceed to the drilling decision model. Details on discretization of prices, information, and the transition matrices of each are available in Appendix B.

Table 3 shows estimates for the full model with all three components: royalty rates, drilling decisions, and production. The signs of coefficients from

[13] Hotz and Miller (1993)-style CCP estimation can accommodate unobserved heterogeneity by using the Expectation-Maximization (EM) algorithm and allowing for a finite mixture distribution as shown by Arcidiacono and Miller (2011). However, I choose to use NFXP estimation since I have a large state space and many not have enough observations to observe all combinations of states multiple times. Additionally, MSL accommodates both the presence of two unobserved variables, ψ_{i0} and ψ_{i1}, plus the additional measurement equations in a very natural way. With the NFXP algorithm, an inner loop solves the firm's value function given a trial guess for the parameter vector, and an outer loop searches for the parameter vector that maximizes the simulated likelihood.

the royalty-rate equation, Equation (1), are as expected. Firms' willingness to pay variables both have positive coefficients, as expected. Their prior, ψ_{i0}, has a very statistically significant impact. This implies that royalty rates are indeed correlated with unobserved heterogeneity in geology and helps rationalize why firms accelerate drilling in high-royalty sections (see Figure 5) and why royalty rates are positively correlated with production outcomes (see OLS estimates in Table 1). While locations with observably better geology (high OGIP) tend to fetch higher royalty rates, this coefficient is surprisingly not significant. The coefficients on landowners' willingness to accept variables also have the expected signs: areas with higher housing values and out-of-state owners require higher royalty payments. Locations with a greater share of permeable surface (less concrete and development) require lower royalty rates.

The primary parameters of interest in firms' drilling problem have to do with well production and firms' information about it that determines output. The two main production parameters—α_g and α_ψ are the coefficients on log OGIP and ψ_1, respectively. So that the model is internally consistent, these are restricted to be the same values in both firms' revenue function and the production equation. This restriction implies very strong identification for these coefficients. The structural estimate of α_g increases to 0.72 compared to the OLS estimate of 0.52 in Table 1. Even though α_g is about one and a half times α_ψ, the standard deviation of log OGIP is only 0.25 compared to a standard deviation of 1.0 for ψ_1. That implies that unobserved heterogeneity (ψ_1) accounts for a much greater proportion of the variation in production than does observable geology. Structural estimates of the production-equation coefficients on time trends and well length are consistent with reduced form estimates in Table 1. The standard deviation of well-specific unobservables, σ_u, is just slightly larger than α_ψ. This implies that even within a section, which is a fairly small area, geological productivity exhibits a fair amount of variation. The last important parameter in the revenue function is ρ, which determines how informative firms' initial signals about geology are. It is the correlation of their prior, ψ_{i0}, and the actual unobserved quality, ψ_{i1}. The estimated value of ρ is 0.66, implying that while fairly informative, firms' initial beliefs are not perfect. There is some informational value to drilling an initial well, and this should serve to accelerate the probability that firms drill an initial well.

Figure 8 breaks down the contribution of the various components of cost over time. The most striking feature of the plot is the large black portion at the top, which is a step function of year effects.[14] This black area represents a regional "shadow cost" of drilling over and above the national drilling cost index. No wells were drilled before the red dashed line. Thus, the very high value for cost before September 2007 is due more to the functional form than variation in the data. Without this shadow cost, the model is unable to reconcile the fact that firms did not drill immediately upon leasing when prices were very high during 2008, but have drilled when prices were much lower in subsequent years. This shadow cost reflects two important considerations. First, the years 2008–2009 coincided with a global financial crisis, and an extraordinary degree of uncertainty in capital markets required to finance drilling. Second, the model is capturing a time in which the oil and gas industry was undergoing a structural shift, not a steady state. The Haynesville was a brand new shale play, and firms' attention was on acquiring mineral leases in the Haynesville at that time. They were not prepared to scale up drilling yet, either because they needed time to build up their supply chains, or because they were still learning about how to adapt their experience in other plays to the Haynesville.

Three other components of the cost function bear comment. First, per-well cost is higher for one well compared to multiple wells ($\alpha_{cheb_{0,1}} < \alpha_{cheb_{0,2+}}$). This reflects the fact that it is costly to move drilling rigs, so average costs per well fall when firms can drill more than one well at a time. Second, the coefficient on a national drilling cost index is negative and significant as expected (i.e., $\alpha_{drilling_t} < 0$). Third, the cash payments required to extend leases ($\alpha_{extension}$) are meaningful and appear to drive firms' drilling decisions.

To assess the fit of the drilling model, I compute the distribution of

[14] More precisely, it is the sum of linear, quadratic, and cubic components of a Chebyshev polynomial transformation of the year.

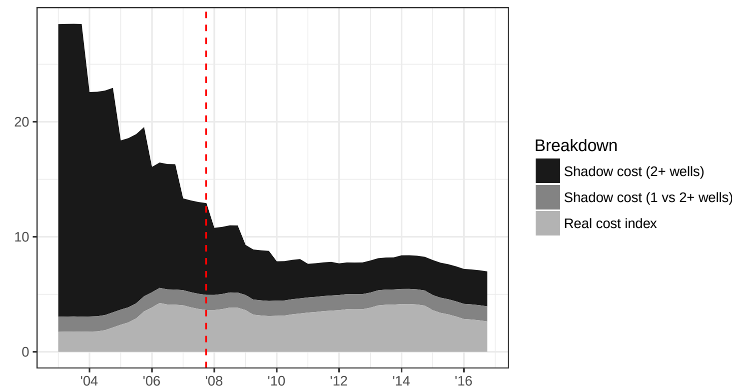

Figure 8: Breakdown of cost components over time

drilling decisions conditional on royalty rates, mineral leases, and exogenous variables. Because royalty rates are both an initial condition as well as correlated with ψ_{i0}, I integrate over the probability of the unobserved heterogeneity given royalty rates, that is, $\Pr(\psi_{i0}, \psi_{i1}|r_i, x_{ri}, g_i)$. This is like a Bayesian posterior distribution for ψ_{i0} and ψ_{i1}. Figure 9a, the top plot, shows the estimated expected number of wells drilled per section in October 2016 (gray) versus the actual (black). This information is also given in Table 10. The model under-predicts the number of sections with 1 or 5–8 wells, and over-predicts the number of sections with 2–4 wells. When α_ψ and α_g are not constrained to be the same in the drilling and production equations, we can better rationalize such a pattern by increasing the role of unobserved heterogeneity, α_ψ. This increases the informational value of well 1 and also increases the likelihood that locations in the upper tail of ψ_1 are totally exhausted. Figure 9b, the bottom plot, shows the expected number of initial and subsequent wells drilled each month. The top pane clearly shows the under-prediction of initial drilling, and the bottom shows where subsequent drilling is over-estimated. The model appears to have some difficulty matching the timing of both initial and subsequent wells—predicted drilling starts too late for Well 1, and starts too early for Wells 2+.

The same procedure used to compute the estimates of expected drilling

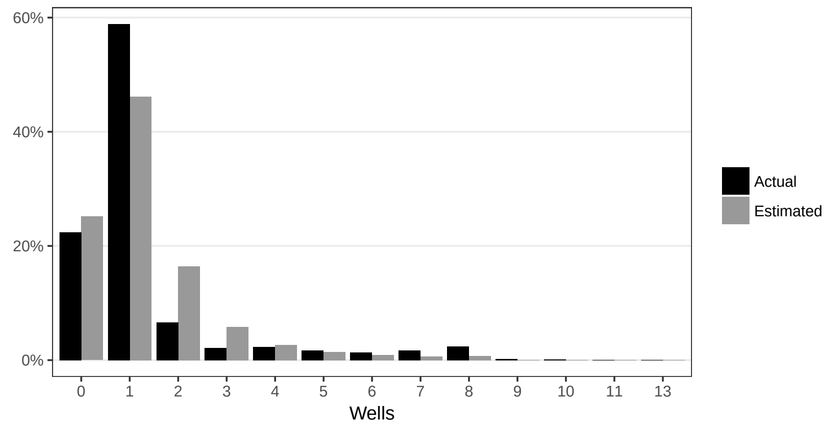

(a) Cumulative drilling in Oct 2016

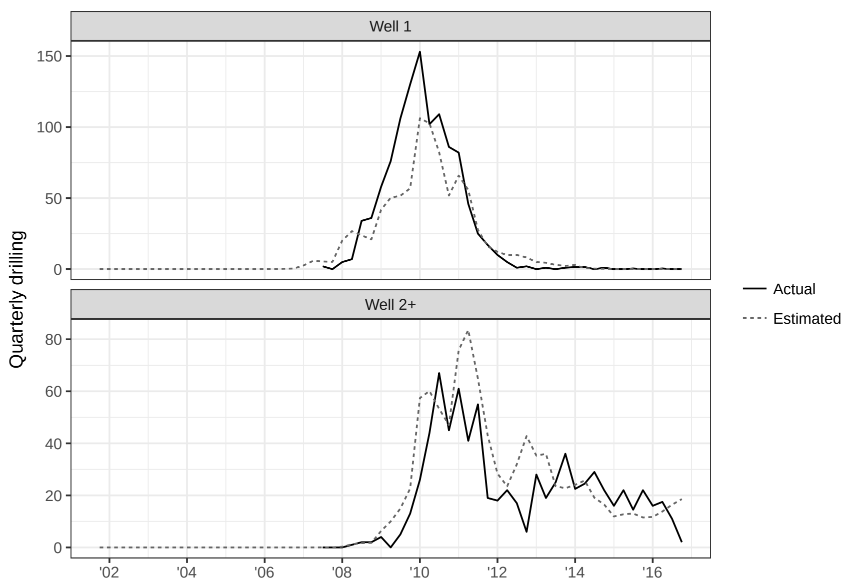

(b) Drilling rate over time

Figure 9: Model fit for drilling decisions

Figure 10: Expected productivity of all wells, initial wells, and later wells drilled

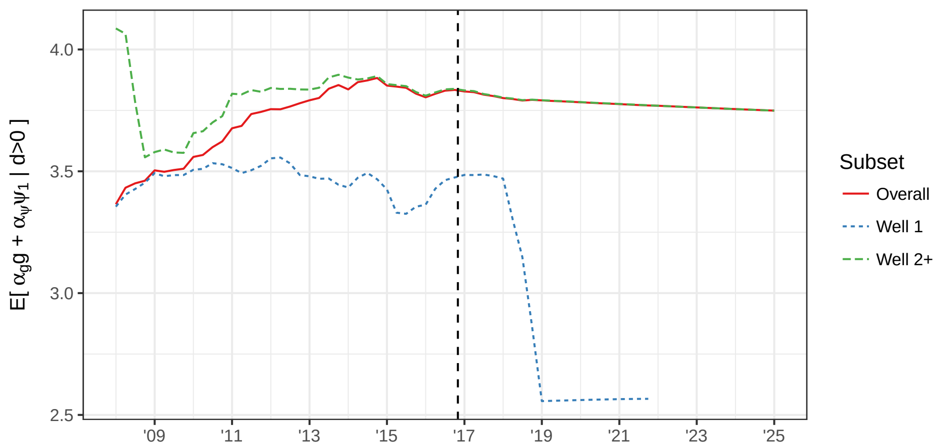

in Figure 9 can be used to compute the expected output of wells drilled over time. Figure 10 shows the expected logarithm of well productivity (times a scalar): that is, $\mathbb{E}[\alpha_\psi \psi_{i1} + \alpha_g \log OGIP | d_{it} > 0]$. We can interpret changes in the level of this quantity as percentage changes in output per well. The blue dotted line represents productivity of Well 1. The green dashed line represents the productivity Wells 2+, and the red solid line, the productivity of all wells. Firms are willing to drill lower-quality sections once to preserve their option to drill and gain information, but they only drill additional wells on higher quality locations. In aggregate, the transition from drilling Well 1 to Wells 2+ causes a rise in mean output per well: the difference from the start of the red line to its peak represents around a 30% increase in output per well. The small variance of $\log OGIP$ implies that the increase in the red line is due primarily to firms drilling locations with larger ψ_1. The right side of Figure 10 after the dashed line shows what happens when we fix prices and costs at their October 2016 values and simulate forward. The implied decline in mean output per well drilled starting after 2016 is due to firms gradually depleting better locations and turning to worse ones. The decline is mild, and firms should be able to overcome it with technology.

Figure 11 illustrates the role that learning about geology plays. The plot shows the difference between expected productivity in a counterfactual, perfect information world in which firms perfectly know the productivity of each section (the top, solid green line with $\rho = 1$) and predicted mean productivity under baseline information (orange dotted line with $\rho = 0.66$) and the counterfactual world in which firms' priors are uninformative (purple dashed line with $\rho = 0$). The three lines differ in both level and slope. Better information (higher ρ) implies that firms don't have to drill poor locations to learn where better ones are. They can concentrate their drilling on the best locations, and they don't allow high-quality leases to expire. This implies a level shift up in mean output per well. There is also a shift in slope, as the red and blue lines (baseline and uninformative worlds) rise compared to the green (perfect information). This increase is the learning effect associated with the fact that firms must drill worse initial wells to learn where the better locations are. Over the course of 2008–2015, learning implies an approximate 15–20% improvement in output per well compared to a perfect information world. The learning effect is a bit larger for the uninformative scenario compared to the estimated value of $\rho = 0.66$. Figure 12 and Table 11 in the Appendix depict the expected cumulative drilling $E[D_{iT}]$ under the three different information scenarios. The primary difference between the three is the number of sections with zero versus one wells drilled. In the perfect information world, more sections are left undrilled, and conditional on being drilled, those sections see a larger number of wells. In contrast, firms drill more initial wells in the uninformative world to learn about the quality of the locations. Since many turn out to have poor locations, however, they do not pursue further drilling.

A final exercise I conduct is to compute a selection correction term: the expected value of ψ_{i1} given the royalty rate and drilling history. I denote the term $\mathbb{E}[\psi_{i1}|royalty, drilling]$. A simple regression of this quantity on the month each corresponding well is drilled yields a highly significant coefficient of 0.11. This suggests that selection alone leads to an 11% per year average increase in output per well—a massive increase over the course of eight years. To correct OLS estimates of monthly well output for the way in which prices,

Figure 11: Deviation of expected productivity compared to perfect information world

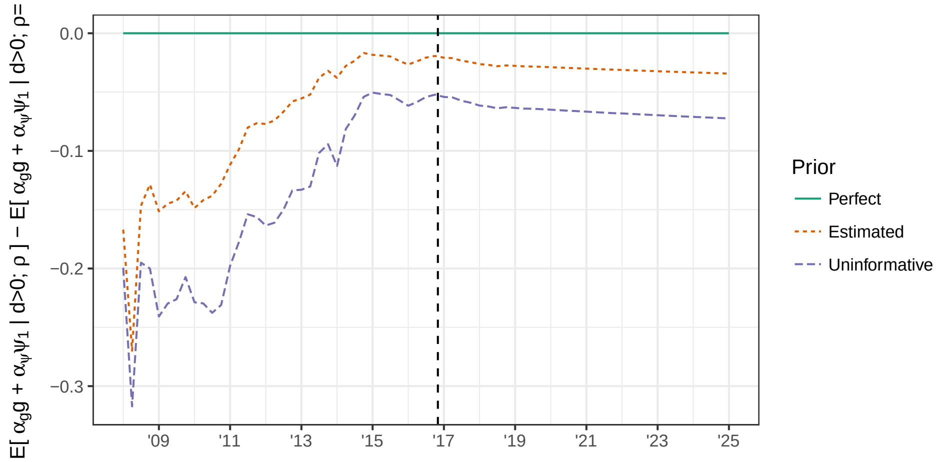

Figure 12: Expected drilling in October 2016 under information settings

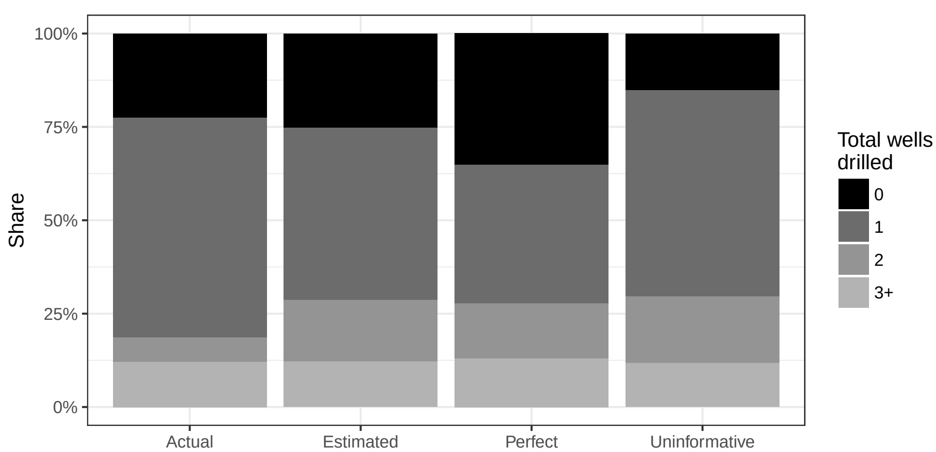

costs, and royalty rates change how firms select particular grades of geology to drill, we can include $\mathbb{E}[\psi_{i1}|royalty, drilling]$ as a regressor.

Table 4 includes uncorrected OLS estimates for monthly well production with and without a subset of (potentially) endogenous variables. The second column (Unrestricted) merely adds $\mathbb{E}[\psi_1|royalty, drilling]$ as a regressor. The coefficients on whether a well is an initial well flips sign and becomes positive and significant, a result that violates what we expect. The effect of the blended royalty rate, while still positive and statistically significant, is diminished by around 30%. The selection term enters positively and with high statistical significance, though it is lower than the α_ψ parameter estimated in the structural estimates listed in Table 3. This suggests that the selection term has not fully corrected for unobserved heterogeneity in geology. Such a result is not entirely surprising given the difficulties that the model has in fitting firms' drilling decisions into such a highly structured framework.

Because the coefficients on royalty rates and whether a well is the first in a section violate theoretical restrictions, the third column (Restricted) zeros out these coefficients. The coefficients on the time trend and selection term both remain positive but exhibit mild, statistically insignificant drops. After replacing two endogenous variables—royalty rates and whether a well is the first to be drilled—the estimated annual increase in productivity, which we might interpret as technology, drops from an estimated 0.49% to 0.42%. Over an 8 year period (roughly the 2008–2016 window of drilling we observe), this implies exogenous improvements in technology have contribute a 34% increase instead of a 39% increase in output per well. Finally, the last column (No time) restricts technology to have no effect. The coefficient on the selection correction jumps by over a standard deviation. We should expect such a result since the passage of time is correlated with an increase in unobserved quality due to the transition from drilling initial wells to drilling subsequent development wells. Technology, like geological heterogeneity, is not observable and may not advance in a linear fashion as time does. Thus, it is difficult to say definitively whether the "Restricted" vs "No time" specification is more correct.

To visually illustrate the importance of correcting for well-specific factors when considering technological improvements over time, Figure 13 plots two variables. The first pane on the left shows the raw, well-level fixed effects from the "Well FE" model in Table 1. This model only controls for deterministic production decline. The second pane on the right shows the residuals from a regression of these raw, well-level fixed effects on log well length, log OGIP, and the selection correction (expected value of ψ_{i1} conditional on royalty rates and drilling histories). Both panes include an estimated time trend. The naive time trend on the left is nearly three times steeper and suggests a 13.8% annual improvement in well performance. The corrected sample on the right suggests a much smaller 4.9% annual improvement. This is still a meaningful number, but *much* more modest.

Figure 13: Well FE and Residuals from Well FE regression

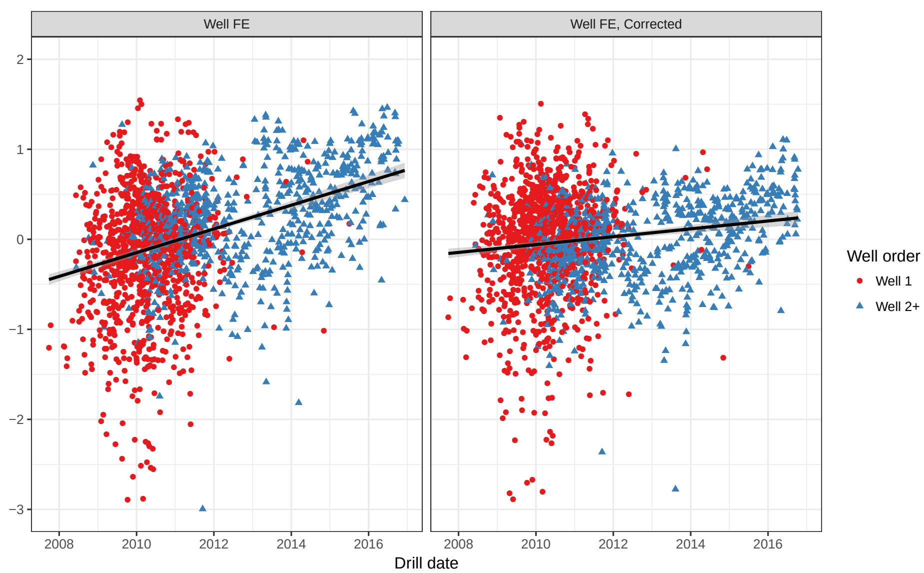

8 Conclusion

It is generally accepted that operators have improved the productivity of their wells by learning how to drill and complete wells better. However, the geological quality of shale formations varies rapidly and widely over space, so *where* firms choose to drill also matters a great deal to output per well. Moreover, price and cost trends, the structure of mineral leases, and the possibility of learning about geological quality all imply that the locations firms select to drill should change over time. To uncover the role that firms' selection of *where* to drill plays in determining output per well, I impose an internally consistent, structural model of royalty rate setting, drilling decisions, and well production. I turn to Louisiana's Haynesville shale for data and exploit the regulatory institutions of there to structure firms' decisions and my data in a consistent, tractable way.

The model reveals that firms' selection of where to drill has important effects on aggregate output per well. The biggest factor in this regard is the shift from drilling initial wells that hold leases by production and provide information about location-specific quality. Simulations suggest that such a shift could rationalize an approximately 30% increase in output per well. Mineral lease expirations provide a powerful incentive to drill worse locations in the hopes that they will become more profitable at higher prices. Second, despite advances in seismic sensing technology, firms lack perfect information about the quality of a particular location. This is consistent with anecdotes about private equity firms that lease acreage, drill initial wells to hold and "prove" the quality of the underlying reservoir, and resell the mineral leases to bigger operators. A comparison of simulations using estimated model parameters with counterfactual simulations in which firms have perfect information suggests that learning about section-specific geology could imply a 15-20% increase in average output per well.

As Figure 13 emphasizes, correcting for well-specific variables, including unobservables, can make a big difference in estimation of exogenous productivity trends. Estimates in Table 4 suggest that correcting for selection on unobservables reduces the estimated effect of exogenous technological progress from 0.49% to 0.42% per year. While the drop is not statistically significant, it does represent an economically meaningful drop of 5.6% over 8 years. For researchers who are interested in estimating productivity trends in shale, using only wells drilled *after* leases are held by production may be a good approach to avoid the impact of mineral lease expirations and firms learning about geology. This, of course, implies that the sample of wells used to estimate production will be drawn from the upper tail of the productivity distribution, but additional geological and engineering information may be useful to assist in understanding the lower end of the productivity distribution. Eliminating initial wells from a sample offers two advantages and focusing only on wells located in leases that are held by production has two other advantages. First, this eliminates the need to control for the expiration dates of mineral leases, which may be difficult to observed. Second, such an approach eliminates much of the initial boom period in a shale. As the breakdown of estimated costs to drill a single well in Figure 8 show, firms may face internal constraints to drilling immediately during the early years of a shale play. These represent unpriced, unobservable shadow costs that dramatically affect firm behavior but are difficult to measure.

Table 3: Full model

Parameter	Estimate	SE	t-statistic	p-value
	Royalty			
ψ_{i0}	0.272***	(0.053)	5.15	0.00
Log median house value	0.500***	(0.080)	6.25	0.00
Out-of-state owners (share)	1.144***	(0.160)	7.15	0.00
Pct totally permeable land	-1.337*	(0.570)	-2.34	0.02
Log OGIP	0.090	(0.086)	1.04	0.30
0.125 | 0.1667	2.810*	(1.138)	2.47	0.01
0.1667 | 0.1875	3.175**	(1.139)	2.79	0.01
0.1875 | 0.2	3.967***	(1.140)	3.48	0.00
0.2 | 0.225	4.908***	(1.140)	4.30	0.00
0.225 | 0.25	5.577***	(1.140)	4.89	0.00
	Drilling			
α_0	-3.498***	(0.233)	-15.03	0.00
α_g	0.715***	(0.050)	14.44	0.00
α_ψ	0.447***	(0.016)	27.41	0.00
$\alpha_{cheb_{0,1}}$	-10.344***	(0.826)	-12.52	0.00
$\alpha_{cheb_{0,2+}}$	-9.029***	(0.821)	-11.00	0.00
α_{cheb_1}	9.907***	(0.672)	14.73	0.00
α_{cheb_2}	-5.207***	(0.403)	-12.91	0.00
α_{cheb_3}	1.293***	(0.197)	6.58	0.00
$\alpha_{drilling_t}$	-0.846***	(0.162)	-5.23	0.00
$\alpha_{extension}$	-1.207***	(0.079)	-15.32	0.00
ρ	0.658***	(0.041)	16.12	0.00
	Production			
Intercept	2.035***	(0.391)	5.21	0.00
$\log \min\{t, t_{int}\}$	-0.517***	(0.005)	-95.18	0.00
$\max\{\log t - \log t_{int}, 0\}$	-1.337***	(0.004)	-363.86	0.00
Log lateral length	0.843***	(0.038)	22.37	0.00
log OGIP	see α_g in *Drilling*			
ψ_{i1}	see α_ψ in *Drilling*			
σ_u	0.463***	(0.011)	42.12	0.00
σ_ϵ	0.546***	(0.001)	453.29	0.00

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$

Log lik. = -97990.2378. Standard errors use approximated Hessian from BFGS.

Table 4: Well production with and without selection correction

	Uncorrected	Corrected		
		Unrestricted	Restricted	No time
$\log \min\{t, t_{int}\}$	-0.532^{***}	-0.533^{***}	-0.532^{***}	-0.529^{***}
	(0.017)	(0.017)	(0.017)	(0.017)
$\max\{\log t - \log t_{int}, 0\}$	-1.325^{***}	-1.324^{***}	-1.324^{***}	-1.340^{***}
	(0.013)	(0.013)	(0.013)	(0.013)
Log well length	0.988^{***}	0.983^{***}	0.988^{***}	1.032^{***}
	(0.120)	(0.115)	(0.117)	(0.119)
Log OGIP	0.523^{***}	0.684^{***}	0.601^{***}	0.637^{***}
	(0.051)	(0.059)	(0.053)	(0.054)
Is 1st well drilled in section	-0.067	0.137^{**}		
	(0.036)	(0.044)		
Spud date (years since 2000)	0.049^{***}	0.052^{***}	0.042^{***}	
	(0.012)	(0.012)	(0.012)	
Blended royalty rate	3.731^{***}	2.600^{***}		
	(0.507)	(0.537)		
$E[\psi_1 \| royalty, drilling]$		0.261^{***}	0.222^{***}	0.269^{***}
		(0.045)	(0.038)	(0.038)
Num. sections (i)	1109	1109	1109	1109
Num. section-wells (iw)	1874	1874	1874	1874
Num. section-well-months (iwt)	$100,982$	$100,982$	$100,982$	$100,982$

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. Standard errors clustered by section. Intercepts omitted.

References

Arcidiacono, Peter and Robert A. Miller (Nov. 1, 2011). "Conditional Choice Probability Estimation of Dynamic Discrete Choice Models With Unobserved Heterogeneity". In: *Econometrica* 79.6, pp. 1823–1867.

Browning, John, Svetlana Ikonnikova, Frank Male, Gürcan Gülen, Katie Smye, Susan Horvath, Carl Grote, Tad Patzek, Eric Potter, and Scott W. Tinker (2015). "Study Forecasts Gradual Haynesville Production Recovery before Final Decline". In: *Oil & Gas Journal* 113.12, pp. 64–71. URL: `http://129.116.232.161/shale/docs/Haynesville%20Shale%20Gas%20Play.pdf` (visited on 01/24/2017).

Covert, Thomas R. (Feb. 22, 2015). *Experiential and Social Learning in Firms: The Case of Hydraulic Fracturing in the Bakken Shale*. SSRN Scholarly Paper ID 2481321. Rochester, NY: Social Science Research Network. URL: `http://papers.ssrn.com/abstract=2481321` (visited on 05/10/2016).

Farmer, Leland E. and Alexis Akira Toda (Oct. 26, 2016). *Discretizing Nonlinear, Non-Gaussian Markov Processes with Exact Conditional Moments*. SSRN Scholarly Paper ID 2585859. Rochester, NY: Social Science Research Network. URL: `https://papers.ssrn.com/abstract=2585859` (visited on 11/08/2016).

Fitzgerald, Timothy (Aug. 2015). "Experiential Gains with a New Technology: An Empirical Investigation of Hydraulic Fracturing". In: *Agricultural and Resource Economics Review* 44.2, pp. 83–105.

Gülen, Gürcan, Svetlana Ikonnikova, John Browning, Katie Smye, and Scott W. Tinker (Oct. 1, 2015). "Production Scenarios for the Haynesville Shale Play". In: *SPE Economics & Management* 7.04, pp. 138–147.

Halton, J. H. (Dec. 1, 1960). "On the Efficiency of Certain Quasi-Random Sequences of Points in Evaluating Multi-Dimensional Integrals". In: *Numerische Mathematik* 2.1, pp. 84–90.

Hendricks, Kenneth and Dan Kovenock (1989). "Asymmetric Information, Information Externalities, and Efficiency: The Case of Oil Exploration". In: *The RAND Journal of Economics* 20.2, p. 164.

Hendricks, Kenneth and Robert H. Porter (June 1, 1996). "The Timing and Incidence of Exploratory Drilling on Offshore Wildcat Tracts". In: *The American Economic Review* 86.3, pp. 388–407.

Herrnstadt, Evan, Ryan Kellogg, and Eric Lewis (Oct. 31, 2018). "Royalties and Deadlines in Oil and Gas Leasing: Theory and Evidence". URL: `https://sites.google.com/site/erickylelewis/home`.

Hitaj, Claudia, Jeremy G. Weber, and Ken Erickson (June 2018). *Ownership of Oil and Gas Rights: Implications for U.S. Farm Income and Wealth.* URL: https://www.ers.usda.gov/publications/pub-details/?pubid=89324 (visited on 10/17/2018).

Hotz, V. Joseph and Robert A. Miller (July 1993). "Conditional Choice Probabilities and the Estimation of Dynamic Models". In: *The Review of Economic Studies* 60.3, p. 497.

Kellogg, Ryan (Oct. 13, 2011). "Learning by Drilling: Interfirm Learning and Relationship Persistence in the Texas Oilpatch". In: *The Quarterly Journal of Economics*, qjr039.

— (June 2014). "The Effect of Uncertainty on Investment: Evidence from Texas Oil Drilling". In: *American Economic Review* 104.6, pp. 1698–1734.

Lane, J. Thomas, Britt A. Freund, and J. Breton McNab (2015). "Ch. 25 Held By Production Leases: When Are They Actually Held?" In: *Energy & Mineral Law Foundation.* Vol. 36, pp. 965–1013. URL: http://www.emlf.org/clientuploads/directory/whitepaper/lane,%20et%20al_15.pdf (visited on 11/11/2017).

Levitt, Clinton J. (2009). "Learning through Oil and Gas Exploration". URL: http://economics.yale.edu/sites/default/files/files/Workshops-Seminars/Industrial-Organization/levitt-100401.pdf (visited on 05/10/2016).

Lin, C.-Y. Cynthia (Dec. 1, 2013). "Strategic Decision-Making with Information and Extraction Externalities: A Structural Model of the MultiStage Investment Timing Game in Offshore Petroleum Production". In: *Review of Economics and Statistics* 95.5, pp. 1601–1621.

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

Montgomery, J. B. and F. M. O'Sullivan (June 1, 2017). "Spatial Variability of Tight Oil Well Productivity and the Impact of Technology". In: *Applied Energy* 195 (Supplement C), pp. 344–355.

Muehlenbachs, Lucija (Feb. 2015). "A Dynamic Model of Cleanup: Estimating Sunk Costs in Oil and Gas Production". In: *International Economic Review* 56.1, pp. 155–185.

Patzek, Tad W., Frank Male, and Michael Marder (Mar. 12, 2013). "Gas Production in the Barnett Shale Obeys a Simple Scaling Theory". In: *Proceedings of the National Academy of Sciences* 110.49, pp. 19731–19736.

Rust, John (1987). “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher”. In: *Econometrica* 55.5, pp. 999–1033.
— (1994). “Chapter 51 Structural Estimation of Markov Decision Processes”. In: ed. by BT - Handbook of Econometrics. Vol. 4. Elsevier, pp. 3081–3143. URL: `http://www.sciencedirect.com/science/article/pii/S1573441205800200` (visited on 09/09/2016).
Seitlheko, Likeleli (Apr. 21, 2016). “Characterizing Production in the Barnett Shale Resource: Essays on Efficiency, Operator Effects and Well Decline”. Thesis. URL: `https://scholarship.rice.edu/handle/1911/96595` (visited on 03/22/2018).
Smith, James L. (2017). “Estimating the Future Supply of Shale Oil: A Bakken Case Study”. In: *MIT Center for Energy and Environmental Policy Research Working Paper CEEPR WP 2017-001*. Massachusetts Institute of Technology.
Smith, James L. and Thomas K. Lee (Sept. 1, 2017). “The Price Elasticity of U.S. Shale Oil Reserves”. In: *Energy Economics* 67 (Supplement C), pp. 121–135.
Steck, Andrew L (Jan. 2018). “Industry Dynamics with Social Learning: Evidence from Hydraulic Fracturing”. URL: `https://sites.duke.edu/asteck/files/2017/11/steck_jmp.pdf`.
Tauchen, George (Jan. 1, 1986). “Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions”. In: *Economics Letters* 20.2, pp. 177–181.
Timmins, Christopher and Ashley Vissing (2014). “Shale Gas Leases: Is Bargaining Efficient and What Are the Implications for Homeowners If It Is Not?” Department of Economics, Duke University. URL: `http://public.econ.duke.edu/~timmins/Timmins_Vissing_11_15.pdf` (visited on 05/10/2016).
Vissing, Ashley (Aug. 2015). “Private Contracts as Regulation: A Study of Private Lease Negotiations Using the Texas Natural Gas Industry”. In: *Agricultural and Resource Economics Review* 44.2, pp. 120–137.
— (2016). “One-to-Many Matching with Complementary Preferences: An Empirical Study of Natural Gas Lease Quality and Market Power”. URL: `http://sites.duke.edu/ashleyvissing/files/2015/10/Full_01092016.pdf`.

A Data details

The DNR website has separate shapefiles for the PLSS grid and the drilling units in the Haynesville. Since not all sections have been unitized, I merge these two datasets. Drilling unit polygons tend to fall on a more regular grid compared to the PLSS sections, so I make some small modifications to the PLSS grid so that it aligns better with the Haynesville drilling units. This is done programatically so as to be replicable.

Of the quarter-million wells in the DNR SONRIS database, 29,458 fall within my geographic definition of the Haynesville, which is taken from (Browning et al. 2015; Gülen et al. 2015). I remove 20,469 wells drilled before January 1, 2000, leaving 8,993 wells to be considered. I define wells to be shale wells if the DNR SONRIS database codes them as a "Haynesville well" (a tax designation) or a horizontal well, or if the well is included in the DNR's "Haynesville wells" shapefile. The Haynesville shale formation and the associated unconventional wells are quite deep, so I further exclude wells shallower than 8700' as well as those drilled into the shallower Fredericksburg or James Lime formations. I also exclude expired permits to drill, injection wells, and abandoned wells as these will not hold leases by production. I exclude several wells that appear to be double-counted or that appear to be associated with one firm targeting the Cotton Valley in a section when another firm is targeting the Haynesville in the same section. Finally, I exclude two dry wells from my sample. Though this introduces a small bias upwards in production estimates, this is small compared to the more than 1000 wells in my final sample, and these dry wells cannot hold leases by production. This leaves 3,619 Haynesville wells that I will consider.

Merging wells to sections involves matching the overlap of units with the line segments that connect wellheads (the location of the vertical part of the well) and bottom-holes (which terminate at the end of the horizontal part of the well). There are no rules for how firms name their wells, but many name them according to the drilling unit names. I also use this information to merge wells and sections. For all but a very few cases, the name and spatial merges concur, and I examine the others on a case-by-case basis. This method of merging is more accurate than using the wellhead location alone since, as Figure 2 shows, the vertical portion of a well may sit in one section when the horizontal wellbore is actually underneath a neighboring section.

I merge production data from commercial provider Drillinginfo to each well based on the well's API number. While the DNR does report production data , it does so at varying levels of aggregation: the lease, unit, or well.

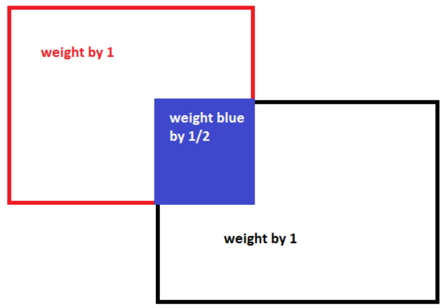

Figure 14: Lease weighting method

Drillinginfo allocates production streams to appropriate wells accounting for whether multiple wells contribute to the same production stream, natural well decline, and well test volumes.

With the mineral leasing information, I keep 68,795 contracts classified by Drillinginfo as a Lease, Lease amendment, Lease extension, or Memo of Lease. I remove 2,434 contracts classified as Assignment, Lease option, Lease ratification, Mineral Deed, Other, or Royalty Deed.

I drop sections for several reasons. The first is that they are on the outer-periphery of the shale and missing geology, or that they are missing median housing values (440 sections). I drop 351 sections that are in areas the Census Bureau classifies as urban in 2010: Shreveport (the large area at the top left) and Mansfield (the smaller one). These urban sections systematically have much higher royalty rates and lower drilling activity than the rest of the sample. Drilling in an urban area is likely to be much more more costly than in a rural one because of more stringent environmental regulations and congestion issues. Moreover, mineral ownership patterns are likely different from those in rural areas. I drop 410 units with "nonstandard" leases that are longer than 10 years or which were signed before January 2003. The vast majority of private leases are less than 10 years long, and the longer leases tend to be on property owned by the government or other large institutions. Because these very long leases between institutions are more likely to have additional requirements, I exclude them. I also exclude the pre-2003 leases, as these pre-date most shale-related activity nation-wide and not likely to be intended for shale development. I drop 78 sections that are smaller than 500 acres or greater than 1000 acres. These primarily occur along the border

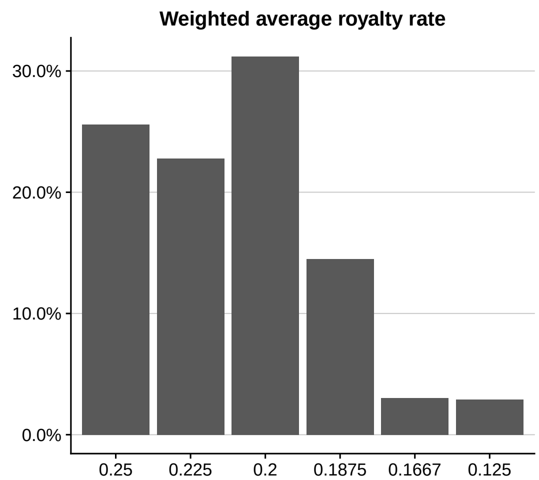

Figure 15: Distribution of discretized, averaged royalty rates r_i (unit-level)

with Texas or in urban areas. The vast majority of lease extension are 24 months: landmen talk about a standard "three year lease with a two year kicker." There are, however, some sections that have leases with extensions that are not 24 months, and I drop them. These sections with nonstandard contracts potentially differ from the others in systematic ways, and handling additional extension lengths requires significantly enlarging the domain of the value function I must compute. For 114 sections, I am missing information on production data (88) or production data and well-length (26). The lack of this information is unlikely to be random: these wells are likely to be conventional or uncompleted. There are 618 sections that I drop where I have leases, but the royalty information is either not present or takes unusually low or high values. The majority of these (453) are dropped for other reasons as well and occur in the rough diagonal from top right to middle left that contains conventional drilling. Similarly, I drop 409 sections that only see conventional drilling as firms are not pursuing unconventional shale development here. Finally, I drop 62 sections where the initial well that would hold them with production spans multiple units (a “cross-unit” well). These wells present two challenges. First, they are likely to have different costs and payoffs compared to single wells. Second, they imply spatial correlation between neighboring sections that I do not model, and it is unclear whether I should treat the multiple sections as a single unit before the initial well is drilled.

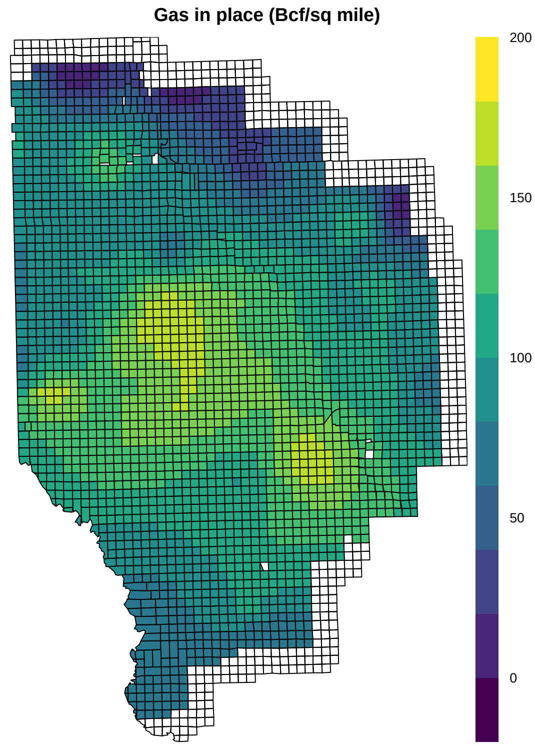

Figure 16: Original gas in place (Gülen et al. 2015)

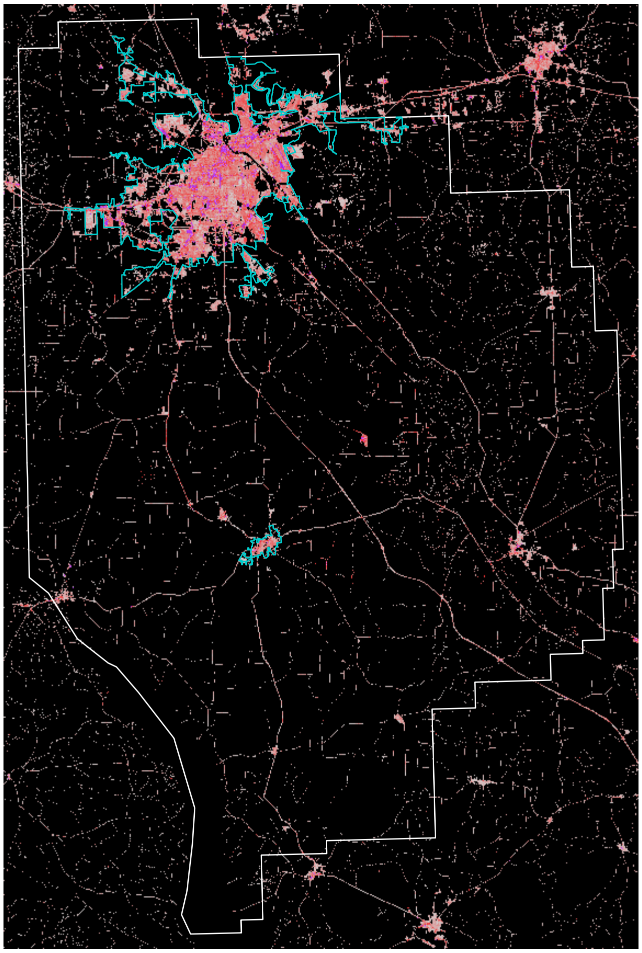

Figure 17: Imperviousness (pink) and urban areas (blue outline)

Table 5: Reasons I drop sections

		Counts			Share		
	Total	Multiple	+ Single	= Either	Multiple	+ Single	= Either
Missing demographics or geology	3158	213	227	440	0.07	0.07	0.14
In Urban area	3158	164	187	351	0.05	0.06	0.11
Nonstandard or missing lease	3158	396	14	410	0.12	0.00	0.13
Section too big (>1000) or small (<500 acres)	3158	59	19	78	0.02	0.01	0.02
Extension not 24 months	3158	200	174	374	0.06	0.06	0.12
Missing production or well data	3158	48	66	114	0.01	0.02	0.04
Nonstandard or missing royalty	3158	453	165	618	0.14	0.05	0.20
Only conventional wells drilled	3158	291	118	409	0.09	0.04	0.13
First well is cross-unit well	3158	42	20	62	0.01	0.01	0.02
Any reason above	3158	754	990	1744	0.24	0.31	0.55

1414 sections remain in sample after dropping for above reasons.
"Multiple" indicates section is dropped for multiple reason; "Single", for one reason only.
"Either" indicates section dropped for multiple or this reason alone.

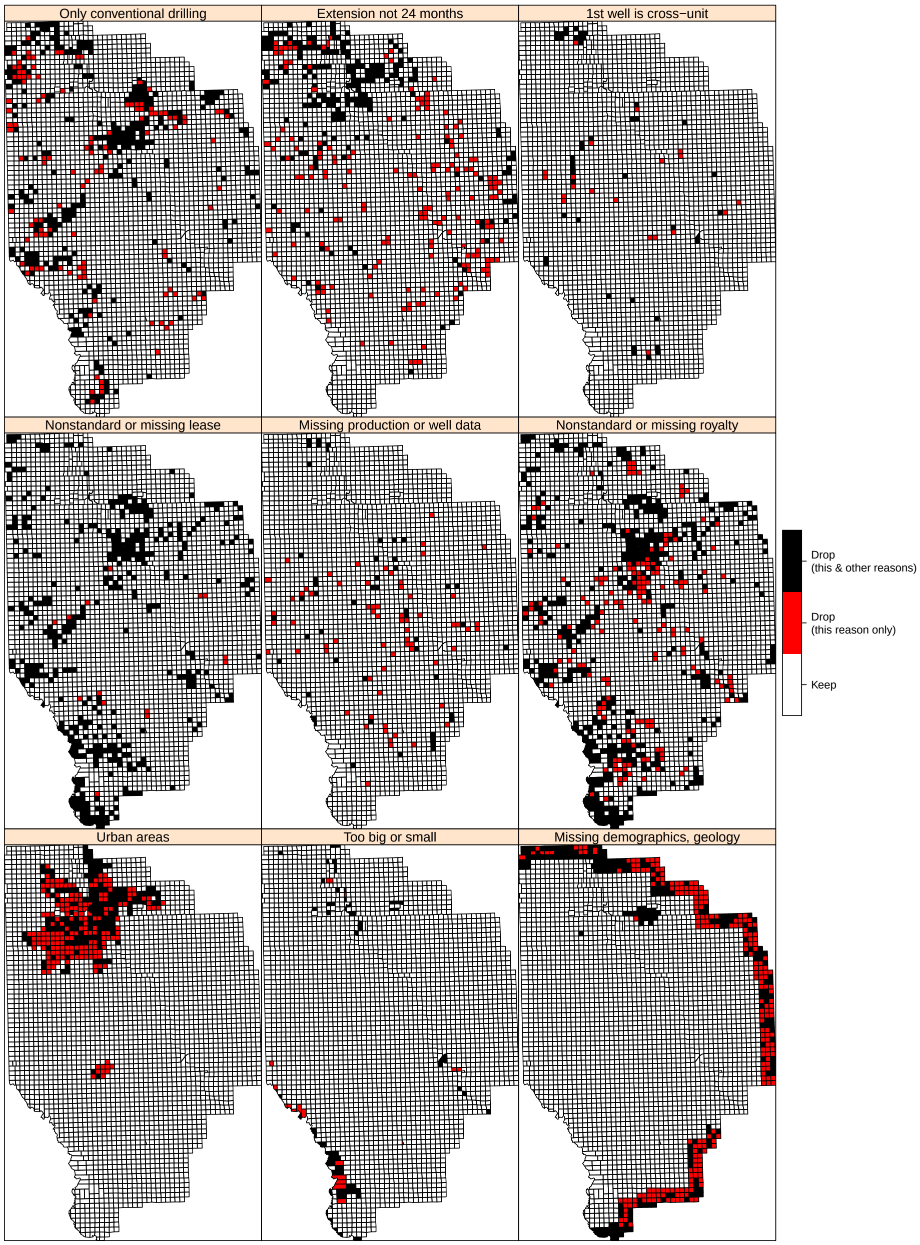

Figure 18: Sections dropped from final sample

Figure 19: Cumulative weekly failure rate by well-order for 36-month leases

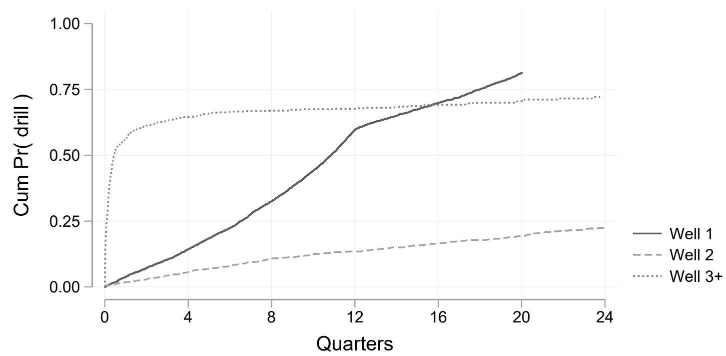

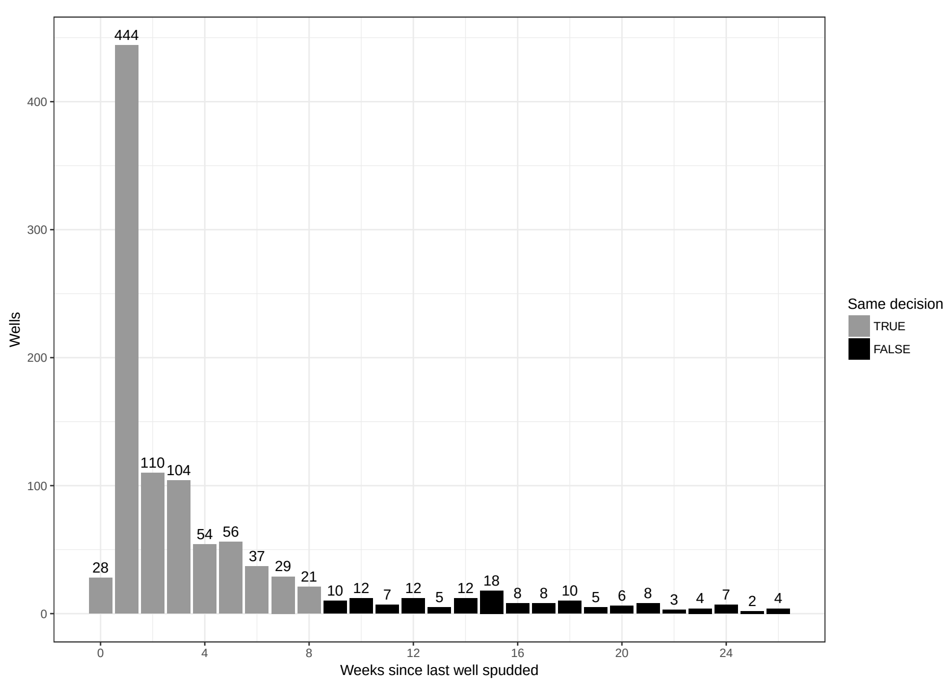

Figure 20: Weeks since previous well drilled

Table 6: Summary: Sections

	N	Mean	SD	Min	Q1	Median	Q3	Max
Acres	1414	645.54	39.33	501.98	635.69	642.91	649.50	962.92
Num shale wells	1414	1.41	1.80	0.00	1.00	1.00	1.00	13.00
0 wells	1414	0.22	0.42	0.00	0.00	0.00	0.00	1.00
1 well	1414	0.59	0.49	0.00	0.00	1.00	1.00	1.00
2+ wells	1414	0.19	0.39	0.00	0.00	0.00	0.00	1.00
Number conventional wells	1414	0.72	2.26	0.00	0.00	0.00	0.00	24.00
First lease signed (year)	1414	2006.72	1.25	2003.50	2005.67	2006.50	2007.75	2014.25
Last lease signed (year)	1414	2009.23	1.50	2003.67	2008.33	2009.25	2010.00	2016.00
Blended royalty rate	1414	0.21	0.03	0.12	0.20	0.20	0.25	0.25
Log OGIP	1414	4.67	0.34	1.34	4.52	4.71	4.89	5.19
Log median housevalue	1414	11.22	0.37	9.79	11.04	11.23	11.38	12.60
Log pop. density	1414	2.08	0.90	0.80	1.37	1.89	2.74	5.39
Share of permeable land	1414	0.96	0.05	0.40	0.94	0.97	0.99	1.00
Share of mineral owners OUT of state	1414	0.10	0.19	0.00	0.00	0.00	0.12	1.00
Share of mineral owners IN of state	1414	0.22	0.27	0.00	0.00	0.11	0.38	1.00
Share of mineral owners with address unkown	1414	0.68	0.33	0.00	0.43	0.77	1.00	1.00
Num leases	1414	17.86	24.26	1.00	4.00	10.00	22.00	355.00

Table 7: Summary: Wells

	N	Mean	SD	Min	Q1	Median	Q3	Max
Log well length	1874	8.35	0.32	2.40	8.31	8.39	8.42	9.08
Log OGIP	1874	4.80	0.25	1.34	4.64	4.82	4.98	5.19
Mean royalty rate	1874	0.21	0.03	0.12	0.20	0.21	0.24	0.26
Num units spanned	1874	1.10	0.32	1.00	1.00	1.00	1.00	3.00
1 unit only	1874	0.90	0.30	0.00	1.00	1.00	1.00	1.00
2 units only	1874	0.09	0.29	0.00	0.00	0.00	0.00	1.00
3 units	1874	0.00	0.07	0.00	0.00	0.00	0.00	1.00
Year drilled	1874	2011.24	1.87	2007.67	2010.00	2010.75	2011.75	2016.92
Year action taken	1874	2011.16	1.87	2007.58	2009.92	2010.67	2011.67	2016.83
Year permitted	1874	2011.10	1.86	2007.58	2009.92	2010.58	2011.58	2016.83
Haynesville well tax designation	1874	0.92	0.27	0.00	1.00	1.00	1.00	1.00
Permitted as cross-unit well	1874	0.09	0.29	0.00	0.00	0.00	0.00	1.00
Log cumulative production	1874	14.96	0.59	10.57	14.70	15.03	15.32	16.20
Months of production	1874	53.89	11.45	1.00	59.00	59.00	59.00	59.00
First month of production data (date)	1874	2011.84	1.91	2008.08	2010.50	2011.33	2012.42	2017.75
First month of production data (month)	1874	2.00	0.00	2.00	2.00	2.00	2.00	2.00
Last month of production data (date)	1874	2016.31	1.35	2009.00	2015.42	2016.25	2017.25	2018.25
Last month of production data (month)	1874	54.89	11.45	2.00	60.00	60.00	60.00	60.00

	N	Mean	SD	Min	Q1	Median	Q3	Max
Time remaining (including extension)	280465	12.16	5.90	0.00	8.00	12.00	17.00	40.00
Observation is during lease extension	280465	0.16	0.36	0.00	0.00	0.00	0.00	1.00
Num wells drilled this month	280465	0.07	0.28	0.00	0.00	0.00	0.00	8.00

Table 8: Summary: Periods

	N	Mean	SD	Min	Q1	Median	Q3	Max
Drilling last period	28757	0.05	0.22	0.00	0.00	0.00	0.00	1.00
Num wells drilled this month	28757	0.03	0.33	0.00	0.00	0.00	0.00	7.00

Table 9: Summary: Leases

	N	Missing	Mean	SD	Min	Q1	Median	Q3	Max
Start (year)	21206	0	2008.23	1.57	2003.50	2007.17	2008.33	2009.42	2016.00
Primary end (year)	21206	0	2011.28	1.60	2006.75	2010.17	2011.33	2012.50	2024.25
Has extension	21206	0	0.79	0.41	0.00	1.00	1.00	1.00	1.00
Extension end (year)	16762	4444	2013.18	1.60	2009.00	2011.92	2013.25	2014.50	2020.83
Primary term (months)	21206	0	36.53	4.88	3.00	36.00	36.00	36.00	120.00
Extension (months)	16762	4444	24.00	0.00	24.00	24.00	24.00	24.00	24.00
Primary + Extension (months)	21206	0	55.50	10.03	3.00	60.00	60.00	60.00	120.00
Has royalty	21206	0	0.22	0.42	0.00	0.00	0.00	0.00	1.00
Royalty	16439	4767	0.22	0.03	0.02	0.19	0.20	0.25	1.75
Royalty < 0.20	16439	4767	0.26	0.44	0.00	0.00	0.00	1.00	1.00
Royalty $= 0.20$	16439	4767	0.31	0.46	0.00	0.00	0.00	1.00	1.00
Royalty $= 0.25$	16439	4767	0.41	0.49	0.00	0.00	0.00	1.00	1.00
Is Lease	21206	0	0.80	0.40	0.00	1.00	1.00	1.00	1.00
Is Memo	21206	0	0.19	0.39	0.00	0.00	0.00	0.00	1.00
Is Other Type	21206	0	0.01	0.10	0.00	0.00	0.00	0.00	1.00
Units per lease	21206	0	1.30	1.62	1.00	1.00	1.00	1.00	132.00
Lease within 1 unit	21206	0	0.83	0.37	0.00	1.00	1.00	1.00	1.00
Lease within 2 units	21206	0	0.12	0.33	0.00	0.00	0.00	0.00	1.00
Log spatially weighted acreage	21206	0	7.95	2.51	1.58	6.11	7.80	9.66	18.25
Log legal acreage	19606	1600	2.70	1.89	-7.10	1.14	3.00	4.11	8.97
Mineral owner is OUT of state	21206	0	0.15	0.36	0.00	0.00	0.00	0.00	1.00
Mineral owner is IN of state	21206	0	0.28	0.45	0.00	0.00	0.00	1.00	1.00
Mineral owner address unkown	21206	0	0.57	0.50	0.00	0.00	1.00	1.00	1.00

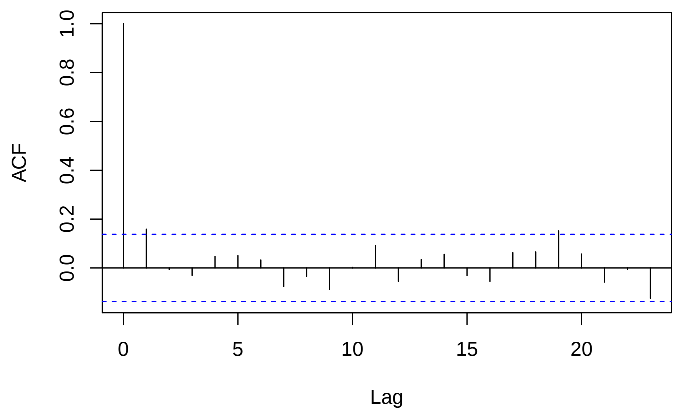

Figure 21: ACF of $\Delta \log p_t$

B Computation

An important element that determines firms' value function is an expectation for the price process. Natural gas prices are generally modeled as having a unit root over the short term, and the ACF of $\Delta \log p_t$, shown in Figure 21, shows very little structure left in $\Delta \log p_t$, suggesting that a random walk is an appropriate way to model $\log p_t$. This does ignore the role of seasonality in natural gas prices. However, the ACF shows no statistical evidence for it, and it is very important to keep the state-space as small as possible for computational feasibility. Given this, I model the price process as $\log p_{t+1} = \log p_t + \omega_{p,t+1}$. The other important process is for drilling costs. I follow Kellogg (2014) and use OLS estimate

$$\log c_{t+1} = \log c_t + 0.0180 - 0.00498 c_t + \omega_{c,t}$$

I assume that ω_p and ω_c are independent normal variables with respective variances 0.100 and 0.050, which I estimate as the standard deviations of $\Delta \log p_t$ and $\Delta \log c_t$. While the latter estimate does differ from OLS, the difference is very slight.

The estimated price volatilities do not enter into a firm's flow-payoffs, but they are extremly important in determining the firm's value function through the transition matrix for prices, Π_z. Since prices are non-stationary, standard discretization methods for stationary AR(1) processes are less helpful. Thus, I use the procedure outlined in Farmer and Toda (2016) to form regime-specific transition matrices, $\Pi_z(k)$, that match as many conditional moments of the continuous distribution as possible. Unfortunately, the joint transition matrix for $\log p_t$, $\log c_t$, and years 2003–2016 is quite large. Thus, I choose a fairly small grid of points: 17, 17, and 14, respectively. While years are discrete, the other two variables are not. Thus, when evaluating $\mathbb{E}V$, I interpolate over $\log p$ and $\log c$ using Quadratic B-Splines. To make sure that boundary conditions for interpolation do not affect interpolated values, I ensure that the grid extends well beyond the range of values we actually observe. Many of the elements of the transition matrices are small and can be ignored with minimal numerical consequences. Thus, I set all elements less than 1e-4 to zero and use sparse matrices. This saves *considerable* computational time.

In addition to discretizing prices, one must also discretize the firm's information about geology, ψ. Since $F(\psi_{i,t+1}|\psi_{it})$ is an AR(1) process, I use the Tauchen (1986) procedure to approximate the transition matrix, Π_ψ.[15] Tauchen (1986) recommends that the grid be symmetric around the mean of the distribution. I set the upper and lower limits to ± 4.0 and use an evenly-spaced grid of 11 points. This covers the vast majority of the distribution as the corresponding quantile is $\Phi(-4.0) = 3.2e-5$. The Tauchen (1986) procedure sets the elements of Π_ψ to be

$$\pi_{ij} = \begin{cases} \Phi\left(x_+\right) & \text{if } j = 1 \\ 1 - \Phi\left(x_-\right) & \text{if } j = 2M_\psi/\Delta_\psi + 1 \\ \Phi\left(x_+\right) - \Phi\left(x_-\right) & \text{otherwise} \end{cases}$$

where

$$x_+ = \frac{\psi_j - \rho^2\psi_i + 0.5\Delta_\psi}{\sqrt{1-\rho}} \qquad\qquad x_- = \frac{\psi_j - \rho^2\psi_i - 0.5\Delta_\psi}{\sqrt{1-\rho}}$$

As with prices, when evaluating the integrated value function, $\mathbb{E}V$, I interpolate between grid-points using quadratic B-splines. This has the added advantage of providing $\partial\,\mathbb{E}V/\partial\psi$ for minimal additional computational cost.

In the inner nested fixed point (NFXP) loop, I solve the integrated value function by backwards induction one leasing-drilling state at a time. The

[15] While the approximation could be improved by either using quadrature grid-points or optimizing over the elements of Π_ψ to match moments, this would require substantial additional computational overhead when computing the inner loop and its derivatives, as well as the likelihood evaluation if the grid changes between iterations. Thus, I choose not to do this.

leasing-drilling state s_{it} is a tuple $s_{it} = (\tau_{0it}, \tau_{1it}, d_{-1,it}, D_{it})$ where τ captures time-to expiration; d_{-1}, whether a well was drilled in the prior period; and D, the cumulative drilling to date. These are sorted lexicographically by $-\tau_1$, $-\tau_0$, $-d_{-1}$, and D. The implication of this is that the integrated value functions at s_i depend on s_j when $i < j$ but not vice versa. The last element in S, $s_{|S|}$, is the the terminal state at which the firm cannot drill, either because the lease expired or all of the possible wells have been drilled. As stated previously, this is normalized to zero: $\mathbb{E}\mathrm{V}(s_{|S|}, z, \psi) = 0 \quad \forall z, \psi$. Computing $\mathbb{E}\mathrm{V}$ at all s involves computing $\mathbb{E}\mathrm{V}$ at $s_{|S|-1}$, then computing $\mathbb{E}\mathrm{V}$ at $s_{|S|-2}$ using $\mathbb{E}\mathrm{V}$ at $s_{|S|-1}$, and so on.

At all leasing-drilling states s_i with $i < |S|$, the firm's problem is finite horizon if the firm cannot remain at s_i by not drilling. Conversely, it is an infinite-horizon problem if the firm can. I solve finite-horizon problems by value function iteration, and infinite horizon problems by a hybrid iteration algorithm that involves a few initial value function iterations and subsequent policy function iterations until convergence (see Rust (1994)). The value function for one set of time-invariant variables, e.g., geology and royalty-rates, does not depend in any way on the value function for another set, which enables parallelization over combinations of the 6 observed royalty rates and 10 geology levels. This considerably accelerates the inner NFXP loops.

The outer NFXP loops involve searching over the simulated likelihoods for a maximum. The log likelihood of each action depends on the flow-payoffs and the integrated value function that correspond to each action in the action space. For each action, I re-compute the flow-payoffs given the state variables and evaluate the value function at the appropriate state values. Because prices, volatility, unobserved information (ψ), and OGIP (G) are continuous state variables, I use quadratic B-splines to interpolate over the value function in these dimensions. I use Monte Carlo integration with two Halton (1960) sequences of bases two and three to integrate out the independent standard normal variables u and v. After discarding the first 5000 observations, for each unit i, I draw 350 pairs of shocks.

I obtain starting values by separately estimating each component of the model and then combining them. Closed-form gradients are available for each component of the likelihood. This allows me to use the BFGS optimization procedure. Conveniently, BFGS stores the inverse Hessian, so I compute standard errors using the BFGS inverse Hessian, as well as the outer product of the gradient. The two are quite close. Estimating the full model is fairly expensive in terms of computational time; however, computation of $\mathbb{E}\mathrm{V}$ and the simulated likelihood are both parallelizable.

C Tables

Table 10: Final cumulative drilling in Oct 2016 (D_T): Actual vs fited

	Count		Share	
D	Actual	Estimated	Actual	Estimated
0	317	356	0.22	0.25
1	832	653	0.59	0.46
2	94	232	0.07	0.16
3	30	82	0.02	0.06
4	33	37	0.02	0.03
5	24	21	0.02	0.01
6	19	13	0.01	0.01
7	24	10	0.02	0.01
8	34	10	0.02	0.01
9	3	0	0.00	0.00
10	2	0	0.00	0.00
11	1	0	0.00	0.00
13	1	0	0.00	0.00

Table 11: Final cumulative drilling in Oct 2016 (D_T): fitted, actual, counterfactual

	Informativeness of prior			
	Uninformative	Estimated	Perfect	
D	$\rho = 0$	$\rho = 0.66$	$\rho = 1$	Actual
0	213	356	497	317
1	781	653	523	832
2	252	232	210	94
3	82	82	83	30
4	35	37	40	33
5	19	21	23	24
6	13	13	15	19
7	9	10	11	24
8	10	10	12	34
9	-	-	-	3
10	-	-	-	2
11	-	-	-	1
13	-	-	-	1

Note: drilling capped to 8 wells in model.

Table 12: Final cumulative drilling in Oct 2016 (D_T) by OGIP (bcf/sq mi): predicted vs actual

	[3.84,92]		(92,111]		(111,133]		(133,179]	
D_T	Actual	Fit	Actual	Fit	Actual	Fit	Actual	Fit
0	153	165	56	94	64	62	44	34
1	174	138	244	171	216	177	198	167
2	19	35	31	55	20	66	24	76
3	4	9	10	17	7	24	9	32
4	1	3	3	7	12	11	17	16
5	2	1	4	3	16	6	2	10
6	0	1	1	2	8	4	10	7
7	0	0	0	1	5	3	19	5
8	1	0	2	1	3	2	28	6
9	0	0	0	0	0	0	3	0
10	0	0	2	0	0	0	0	0
11	0	0	0	0	1	0	0	0
13	0	0	0	0	1	0	0	0

Note: drilling capped to 8 wells in model.

Table 13: Final cumulative drilling in Oct 2016 (D_T) by Royalty: predicted vs actual

	12.5%		16.67%		18.75%		20%		22.5%		25%	
D_T	Actual	Fit	Actual	Fit	Actual	Fit	Actual	Fit	Actual	Fit	Actual	Fit
0	6	16	12	14	53	53	112	113	62	79	72	83
1	28	17	23	19	115	97	246	200	191	149	229	169
2	5	5	4	6	13	33	28	72	26	54	18	62
3	0	2	0	2	6	11	7	26	11	19	6	22
4	0	1	2	1	3	5	12	12	8	9	8	10
5	0	0	0	0	0	3	10	7	5	5	9	6
6	0	0	0	0	3	2	4	4	7	3	5	4
7	0	0	1	0	6	1	7	3	3	2	7	3
8	0	0	0	0	5	1	13	3	8	2	8	3
9	0	0	1	0	0	0	1	0	1	0	0	0
10	2	0	0	0	0	0	0	0	0	0	0	0
11	0	0	0	0	1	0	0	0	0	0	0	0
13	0	0	0	0	0	0	1	0	0	0	0	0

Note: drilling capped to 8 wells in model.