WORKING PAPER

Supply Flexibility in the Shale Patch: Evidence from North Dakota

NORGES BANK RESEARCH

9 | 2017

HILDE C. BJØRNLAND,
FRODE MARTIN NORDVIK
AND
MAXIMILIAN ROHRER

NORGES BANK

Working papers fra Norges Bank, fra 1992/1 til 2009/2 kan bestilles over e-post:
FacilityServices@norges-bank.no

Fra 1999 og senere er publikasjonene tilgjengelige på www.norges-bank.no

Working papers inneholder forskningsarbeider og utredninger som vanligvis ikke har fått sin endelige form. Hensikten er blant annet at forfatteren kan motta kommentarer fra kolleger og andre interesserte. Synspunkter og konklusjoner i arbeidene står for forfatternes regning.

Working papers from Norges Bank, from 1992/1 to 2009/2 can be ordered by e-mail
FacilityServices@norges-bank.no

Working papers from 1999 onwards are available on www.norges-bank.no

Norges Bank's working papers present research projects and reports (not usually in their final form) and are intended inter alia to enable the author to benefit from the comments of colleagues and other interested parties. Views and conclusions expressed in working papers are the responsibility of the authors alone.

ISSN 1502-819-0 (online)
ISBN 978-82-7553-982-1 (online)

Supply Flexibility in the Shale Patch: Evidence from North Dakota*

Hilde C. Bjørnland† Frode Martin Nordvik‡ Maximilian Rohrer§

May 30, 2017

Abstract

We analyse if supply flexibility in oil production depends on the extraction technology. In particular, we ask to what extent shale oil producers respond to price incentives by changing completion of new wells as well as oil production from completed wells. Using a novel well-level monthly production data set covering more than 15,000 crude oil wells in North Dakota, we find large differences in response between conventional and unconventional (shale) extraction technology: While shale oil wells respond significantly to spot future spreads by changing both well completion and crude oil production, conventional wells do not. Our results suggest that firms using shale oil technology are more flexible in allocating output intertemporally. We interpret such output pattern of shale oil wells to be consistent with the Hotelling theory of optimal extraction.

JEL-codes: C33, L71, Q31, Q40

Keywords: Oil extraction, crude oil prices, US oil shale boom, Hotelling theory

*This Working Paper should not be reported as representing the views of Norges Bank. The views expressed are those of the authors and do not necessarily reflect those of Norges Bank. The authors would like to thank Jon Fiva, Lutz Kilian, Marco J. Lombardo, Klaus Mohn, Hyungsik Roger Moon, Gisle J. Natvik, M. Hashem Pesaran, Francesco Ravazzolo, Barbara Rossi, Jim Stock, Leif Anders Thorsrud and Robert J. Vigfusson, as well as conference participants at 9th International Conference on Computational and Financial Econometrics (CFE 2015) in London, the Oil, Middle East and the Global Economy conference 2016 in LA, the 69th European Meeting of the Econometric Society (ESEM 2016) in Geneva, the CAMP 2016 workshop on Commodity prices and Monetary Policy in Oslo, the International Association for Energy Economics annual conference (IAEE 2016) in Bergen and the Forecasting in Finance and Macroeconomics conference at the Free University of Bozen/Bolzano for useful comments and fruitful discussions. This paper is part of the research activities at the Centre for Applied Macro and Petroleum economics (CAMP) at the BI Norwegian Business School. The usual disclaimers apply.

†BI Norwegian Business School and Norges Bank. Email: hilde.c.bjornland@bi.no

‡Norges Bank and BI Norwegian Business School. Email: frode.m.nordvik@bi.no

§BI Norwegian Business School. Email: maximilian.rohrer@bi.no

1 Introduction

Since the summer of 2014, the global oil industry has been in one of its deepest downturns in decades. The cause is the plunging oil price, which at one point fell more than 70 percent compared with the observed price level in June 2014. These events have spurred renewed interest in the behaviour of oil producers, in particular to what extent producers respond to these negative price incentives by reducing oil production and drilling activity.

Previous studies addressing this issue typically find oil production to be price inelastic in the short run, see e.g. Hogan (1989), Pesaran (1990), Dahl and Yücel (1991), Ramcharran (2002), Smith (2009), Griffin and Teece (2016) and Anderson et al. (2014). However, these studies analyse output responses from conventional oil pools only, mainly using aggregate data. Yet, aggregate elasticity of oil supply depends on the extraction technology of the marginal producer of crude oil. Arguably, the present marginal producers are US shale (unconventional) oil firms.[1] In fact, by 2015, oil production from shale deposits accounted for half of total US crude oil output. Furthermore, the use of hydraulic fracturing technology is spreading to other oil-producing countries, potentially making unconventional oil a much larger share of total production than it is today (see Clerici and Alimonti (2015)).

Hence, knowledge of supply elasticity and producer behaviour in unconventional oil pools is important, but lacking. We contribute to the literature by examining the response of shale and conventional crude oil wells to spot and expected future oil prices. Using a novel and rich monthly panel data set from 1986 to 2015, covering more than 15,000 oil wells operated by more than 450 oil-producing firms, both conventional and shale in the North Dakota oil patch, we are able to study the response of crude oil wells along two margins. First, along the *intensive* margin, we ask if firms adjust the flow rate from existing wells in response to oil price shocks, distinguishing between conventional and shale oil wells. Second, along the *extensive* margin, we ask if firms optimise the completion of new wells in response to oil price changes, comparing again the two types of well technologies. The decision to complete a well is equivalent to exercising the real option to produce, since this is the actual start of production from a well. Note that the timing of when to produce is different from the timing of drilling the well. The reason is that wells are frequently drilled, but left uncompleted for a some time. Once a well is completed, however, it starts producing. We postulate that shale firms more often than conventional firms leave wells that are drilled uncompleted, or in below ground storage, so as to optimize the timing of well completion. Shale oil reserves are assets than can be quickly liquidated when markets become favorable.

There are several reasons why one would expect shale wells to respond differently from conventional oil producers. First, shale oil wells have a much higher marginal cost per barrel, increasing the risk that prices enter the range of marginal costs inducing firms to

[1] In the paper, we use the terms 'shale' and 'unconventional' interchangeably.

reduce output or shut in the well temporarily. Second, shale technology is a more flexible well extraction technology, making firms less constrained in taking advantage of real option values. For example, a shale well may be stimulated with water and chemicals many times during its lifetime, sometime as much as 20 times, see for instance Carpenter (2014). This leaves the well operator some degree of freedom to choose the timing of the fracturing operations relative to a conventional well which is naturally flowing. Third, shale wells have a very front-loaded production profile relative to conventional wells. In our sample, the first-month production rate of the average shale well is almost four times as high as that of a conventional well. Moreover, the decline rate of shale wells is on average about five times higher than for conventional wells. This production front-loading increases the incentive for shale producers to optimise the timing of well completion.

Using the full panel of output per well, and distinguishing between the two different well technologies, we first show that neither conventional oil wells nor shale wells respond significantly along either the intensive or extensive margin to changes in spot oil prices, echoing previous results in the literature. However, when we also include expected future oil prices into the analysis, the picture changes dramatically. In particular, we document that (i) shale wells have a positive and significant short-term supply elasticity in response to shifts in the one-month spot future spread. The immediate response is to increase the monthly well flow in shale wells by nearly two percent for a ten percent increase in the spot future spread. Furthermore, shale wells (ii) reduce both well flow rates (intensive margin) as well as the number of new completed wells (extensive margin) they add to the stock of producing wells substantially today when the long term (three-, six-, twelve- or 18-month) future curve is upward sloping (and vice versa when the future curve is downward sloping). The magnitudes are of economic importance. For a ten percent increase in the twelve-month spot future spread, producers respond by reducing the number of wells they complete by 2-3 percent, or on average 8-10 wells pr month. This is equivalent to holding back about 80,000 barrels of monthly crude oil production from the market. We find no evidence of such dynamics for conventional oil producers.

The use of microeconomic data to infer the price elasticity of aggregate output has several advantages. First, by constructing a rich panel data set, we can eliminate any potential aggregation bias over well production rates when estimating the empirical model. Aggregating over all individual wells is equivalent to imposing identical parameter values for all producing wells regardless of well or firm characteristics. This would lead to a significant loss of information about micro relations in our data set. Second, the use of panel data enables us to explore the cross-sectional variation in for instance well type, age, location or other characteristics of interest, and we can investigate the potential heterogeneity in producer behaviour across technologies (conventional or unconventional). Third, using the large cross-section of the panel, we can identify differential behaviour of conventional and shale producers in response to the same price shock. Hence, our results are immune to unobserved time variation in state variables, such as changing market conditions or oil price regimes. Lastly, having a large cross-section in a panel is beneficial for statistical inference when analysing a relatively short time period as we do here.

We interpret the results for shale oil producers to be in line with Hotelling (1931)'s model of optimal exhaustible resource extraction. Reserves are an inventory, and the decision to produce is an intertemporal choice of when to draw down below-ground inventory. For producers to behave in line with the Hotelling (1931) theory, they must be able to reallocate extraction across different periods. Previous studies for conventional producers have not found this flexibility condition to hold, see e.g. Anderson et al. (2014) that found the price elasticity of oil suppliers in Texas to be zero. Our empirical results, however, show that the degree of output flexibility depends on the technology applied, and that firms using shale oil technology are much more flexible in allocating output intertemporally. This enables shale producers to behave more consistently with the benchmark theories for commodity producer behaviour.

Our result has important policy implications. For instance, when designing tax policies, policymakers should take into account that producers adjust differently to price-sensitive news (such as a tax on revenue). Hence, any tax policy that targets operators at the well level should consider that shale producers are more responsive to price changes and can optimise their production profile than what conventional well operators can do. Further, the results have also far fetching implications for oil prices. If marginal supply is shale producers, then we would expect a stabilizing effect on prices. Finally, our results have also important implications for how one should analyse the role of oil in the macroeconomy. Oil price-macro models have often assumed aggregate oil production to be price inelastic in the short run when identifying oil market shocks, c.f. Kilian (2009). However, as production from drilled shale wells will be responsive to shocks to the oil price also in the short term, this assumption may no longer hold. Instead, our results support exploring alternative identification schemes that relaxes the assumption of zero short-run response in oil production to price signals, see for instance Kilian and Murphy (2012) and Lippi and Nobili (2012) that assumes low, but not perfectly inelastic short-run supply curves, or more recently, Baumeister and Hamilton (2015) and Caldara et al. (2016) that allow the short-run supply elasticity to increase further.[2]

The rest of the paper proceeds as follows: Section 2 describes the data environment, while Section 3 sets up the panel data model and discusses the empirical results. Section 4 discusses robustness and some caveats, while Section 5 interprets the results in light of the Hotelling (1931)'s model of optimal exhaustible resource extraction. Concluding remarks follow in Section 6.

[2]For instance, Caldara et al. (2016) proposed an identification scheme that restricts the elasticities by minimizing the distance between the elasticities allowed by the SVAR and some selected target values constructed from a survey of relevant studies (i.e., Hamilton (2009)).

2 Data environment

Below we provide a short background on the methods of oil extraction, we explain how we identify conventional and unconventional wells and finally describe some key features of the data set.

2.1 Oil and geology in North Dakota

The first commercial oil discovery in North Dakota was made in 1951 in the Williston basin. Discoveries continued to be made in the 1960s and 1970s, and production gradually increased until it peaked during the first conventional oil boom in the early 1980s. At the peak (in 1981), 834 conventional wells were drilled in one year. From the peak followed a long period of decline, until only 34 wells were drilled in 1999, the lowest number of wells drilled since oil drilling began. By April 2004, production in North Dakota had also reached a minimum, at around 75,000 barrels produced per day.[3]

New discoveries from 2006 initiated the recent North-Dakota boom. The boom relates mainly to oil deposits discovered and produced within the Bakken field, which is a rock formation occupying an area of about 520,000 km^2, about the size of Spain.[4] The boom started with the discovery of Parshall Oil Field in 2006, and peaked in 2012, but with substantially less growth noted since 2015, see Figure 1. The boom was made possible by the continued development of new extraction techniques. While these techniques had existed for some time, the cost of extraction had been too high for it to be profitable. The increase in oil prices throughout the 2000s made drilling of shale oil competitive with conventional techniques, and investment in unconventional oil started to boom.

The method of oil extraction will depend on the geology of the site. *Conventional* oil well technology requires oil- and gas- bearing rock that is porous, such as sandstone or washed out limestone. When crude oil forms in a permeable[5] rock, gas will naturally gather at the top of the reservoir and the crude oil will be trapped in the porous rock underneath the gas cap. At the very bottom, there is water. The crude oil is pressurised in the pores of the formation rock, so that when a well is drilled into the reservoir, the formation pressure and the permeability of the rock naturally drives the hydrocarbons out of the rock and up into the well. Conventional oil well technology can produce effectively from these types of reservoirs, aided by the natural pressure and the permeability in the reservoir. In practice, this involves drilling a conventional vertical well straight into the reservoir and producing the oil that flows by itself into the well, see Devold (2013) for further details. The law of physics governing the flow of fluid from two locations in

[3]For details, see North Dakota Studies (2016).

[4]According to the US Geological Survey (USGS), see Gaswirth et al. (2013), it is the largest continuous oil resource in the lower 48 states. In april 2013, the USGS estimated that the amount of oil that could be economically recovered from the Bakken would be at 7.4 billion barrels. The formation stretches out over regions such as North Dakota, Montana, Saskatchewan and Manitoba.

[5]Permeability is a geological term to describe how easily oil flows naturally through rock.

Figure 1. New wells over time.

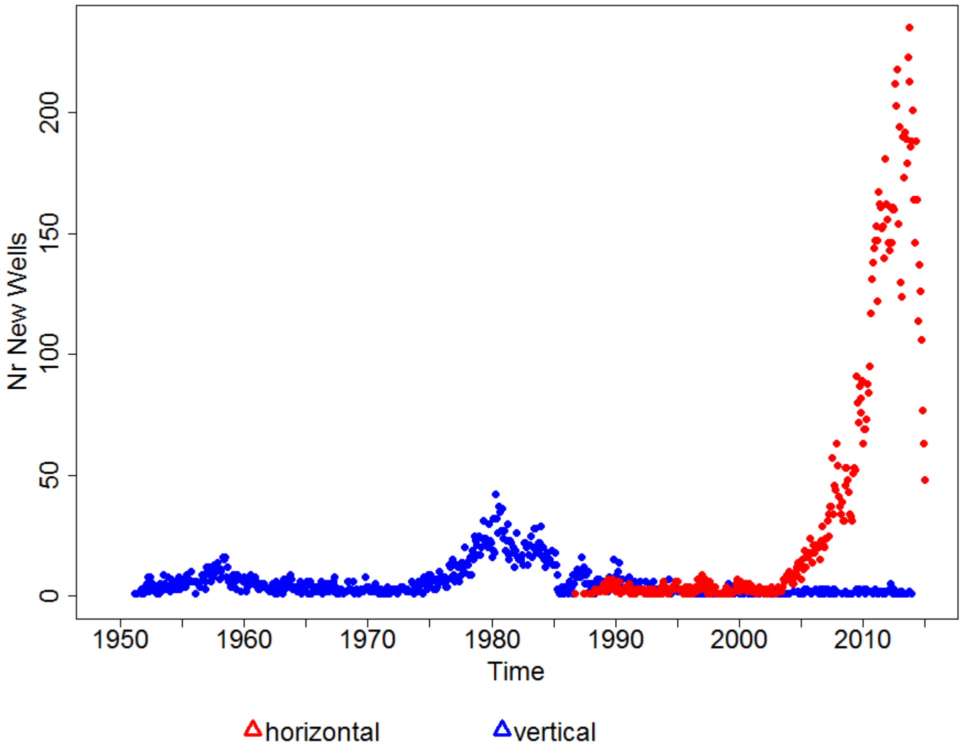

Note: The number of new wells entered into production on a monthly frequency in North Dakota since 1952, separated into well technology type. The blue dots represent conventional/vertical wells, and the red represent shale/horizontal wells. See Section 2.2 for more details on the data.

a porous medium with different pressure is known as Darcy's law, c.f., Hubbert et al. (1956).

In contrast to a conventional oil reservoir, when crude oil is trapped in a rock formation that has zero permeability, the natural pressure in the reservoir formation is not enough to make the oil flow into the well once a well is drilled - because the oil is trapped in small pockets inside the shale rock formation. Naturally, Darcy's law of fluid flow also applies to shale rock, but the zero permeability restricts the flow of fluid.[6] These types of reservoirs, often called tight oil reservoirs, require additional stimulation once the well has been drilled. That is, a combination of hydraulic fracturing (or "fracking") combined with horizontal drilling enables the oil to escape the rock formation.

For both conventional vertical wells and unconventional wells, however, the process of drilling a well starts in a similar way: a rig drills a vertical well into the ground, to depths of up to 10,000 feet. For conventional vertical wells, the rig will stop drilling at this point, and the well will be completed for production as a vertical well. In shale oil wells, however, the well will be drilled further, but at an angle, creating a so-called "bend". This initiates the horizontal part of the well, which can extend up to 10,000 feet in the

[6]See for instance oil field glossary at http://www.glossary.oilfield.slb.com/Terms/p/permeability.aspx for details.

horizontal direction. Once the well is drilled, it is encased with a metal pipe, a so-called casing, throughout the well. This is done to keep the formation wall from caving into the well bore and to be able to control the flow of oil to the well head. For shale wells, the drilling rig and its crew leave the site at this point. The well must then be completed, i.e. fracked by a fracking crew, in order to actually start producing. The fracking crew arrives with a so-called "missile", a manifold that is placed on the well head to pump the fracking water and chemicals down into the well at high pressure to crack the formation open.[7] When the formation is open, the oil starts to flow up to the well head at high pressure. The fracking process can take as little as 2-3 days to be completed, and may be repeated several times during the lifetime of the well.[8] Figure A.1 in Appendix A illustrates this. It shows production profiles for shale wells that have been fracked several times. When a well has been fracked, the released shale oil causes a temporary boom in the production rate, which declines as the effect of the fracturing dwindles. If this process is reiterated, it generates a spiky production profile. This also explains why the average marginal cost of shale oil production is substantially higher than for conventional production.[9]

A particular feature of oil extraction in North Dakota is that, while the conventional oil boom that peaked in the 1980s was a *vertical* drilling boom, the recent unconventional oil boom is a *horizontal* drilling boom, see e.g. Miller et al. (2008) for details. The main reason for this relates to the thickness of the middle Bakken geological layer, which is where most of the shale oil is extracted. The geological layer is at most 150 feet thick, which is fairly thin for an oil producing zone, see e.g. Meissner (1984). This makes it highly inefficient to produce by vertical wells, hence the well is drilled horizontally so more of the wellbore can be exposed to the oil producing zone.

For both well technology types, the feasibility of varying well flow rates varies over the lifetime of the well. In the initial phase, the oil will most often flow naturally to the surface, aided by the pressure difference between the well head and the reservoir. The most productive wells will produce naturally, without any aid, for years. In this phase, varying the flow rate of the well is very costly, and rarely done. However, after this initial phase, the well often needs some sort of artificial pumping unit to bring the oil to the surface in sufficient quantities to make the well cover variable costs. Since shale wells have a decline rate that is very high relative to conventional wells, the phase when the well needs some sort of artificial lifting unit happens at a younger well age, relative to the

[7]Several large trucks with the enginge power to maintain the high pressure are needed in this process, and in addition a blender truck that blends the chemicals, sand and water together. The blender truck then sends the water mix to the missile, which sends the mix down the well at very high pressure in order to crack the tight oil reservoir formation open.

[8]For a short introduction, see Dunn (2016).

[9]Although marginal cost per barrel varies across wells, and over time with varying oil prices, the average marginal cost of shale oil production is estimated at about USD 60-70 per barrel. This is substantially higher than for conventional production, even in deep-water offshore fields, see Gopinath (2014).

conventional vertical wells. Hence, this increases the incentive to optimize the timing of well completion.

2.2 Identifying conventional and unconventional wells

The data set was retrieved from the database of the North Dakota Industrial Commission (NDIC), Oil and Gas Division. It provides production figures on a monthly frequency for 16,639 crude oil wells in North Dakota.[10] The total data set was collected for the full time period from 1952 until 2015. However, the sample period used in the analysis is limited by the availability of future oil price series, and thus ranges from July 1986 to November 2015.

An important issue is how we can distinguish a shale well from a conventional well in the database. As mentioned in Section 2.1, tight oil in North Dakota is found in thin geological layers which can only be efficiently extracted with horizontal well technology (combined with hydraulic fracturing). Hence, the shale oil boom is a horizontal drilling boom, while the conventional oil boom that peaked in the 1980s was a vertical drilling boom. In the data we observe the drilling technology (horizontal vs. vertical) at well level and can thus identify shale oil production by the drilling technology.

We limit our sample to active crude oil wells, but we also allow for one month with zero production (shut in production) whenever this month is preceded or succeeded by a month of positive output. The main advantage of this is that we exclude long time periods of wells being left idle, but we retain the periods when wells are temporarily shut-in. Since it is well known that some well operators shut in their wells temporarily when oil prices are low, we believe that including these wells will yield better estimates of true supply elasticity. Figure 1 displayed in Section 2.1 above graphed the complete history of monthly well entries in North Dakota dating back to the first well which entered into production in 1952. As was seen in the figure, the recent boom in the number of wells entering into production surpasses anything previously seen in North Dakota. Furthermore, the boom is entirely caused by horizontal drilling and fracturing. Figure 2 further details all of North Dakota's crude oil wells from each well's geographic coordinates within the state. We note that the density of wells is high, particularly in certain areas. The shale wells appear to be drilled more closely together than conventional wells, often in an array-like pattern.

From the constructed data set, we observe that two-thirds of all wells are horizontal shale wells. The average age of a horizontal well is 42 months, whereas the average age of a conventional well is 252 months. This underlines a need to control for well age in the analysis, see Table A.1 in Appendix A for main summary statistics. Overall, there are 489 operator firms that operated wells in the sample period, reflecting an industry structure with many small and independent firms. The average number of wells completed

[10] The raw data covers more than 33,000 wells. However, after requiring that there is sufficient data observations, it reduces to 16,639.

Figure 2. North Dakota: Individual Wells

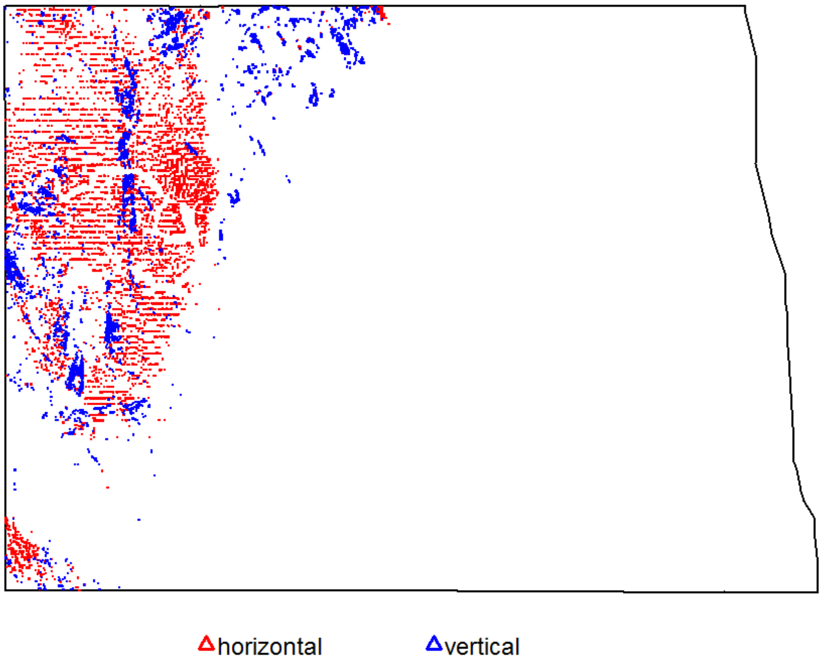

Note: Map of North Dakota plotting each individual oil well's location since 1952 according to latitude and longitude. The blue dots represent conventional/vertical wells, and the red represent shale/horizontal wells.

per firm in the sample period is 20 wells. However, the number of completions is highly concentrated. More than 60 percent of firms only completed one to three wells during the sample period, while a small fraction of firms completed more than 500 wells during the sample period.

2.3 Well drilling, completion and production phase

In this paper, we study both the intensity of production, and the timing of the decision to start to produce, which is the time when the well is completed. As alluded to above, the timing of when to *produce* is different from the timing of *drilling* the well. The reason is that wells are frequently drilled, but left uncompleted for a some time. Once a well is completed, however, it starts producing.

Since the different types of geology in shale and conventional oil reservoirs require different extraction technology, the production profile for shale wells will differ from that of conventional wells. In particular, the shale fracturing process and the high pressure at which oil flows out of the fractured shale rock, suggests much higher initial production rates from shale oil relative to conventional oil. Figure 3 illustrates this. The figure graphs the average monthly production rates for the two types of well technology, plotted against well age. Table 1 displays some corresponding comparative statics. Two main features stand out. First, the production profile of shale wells is heavily front-loaded relative to

Figure 3. Production profiles of horizontal and vertically drilled wells

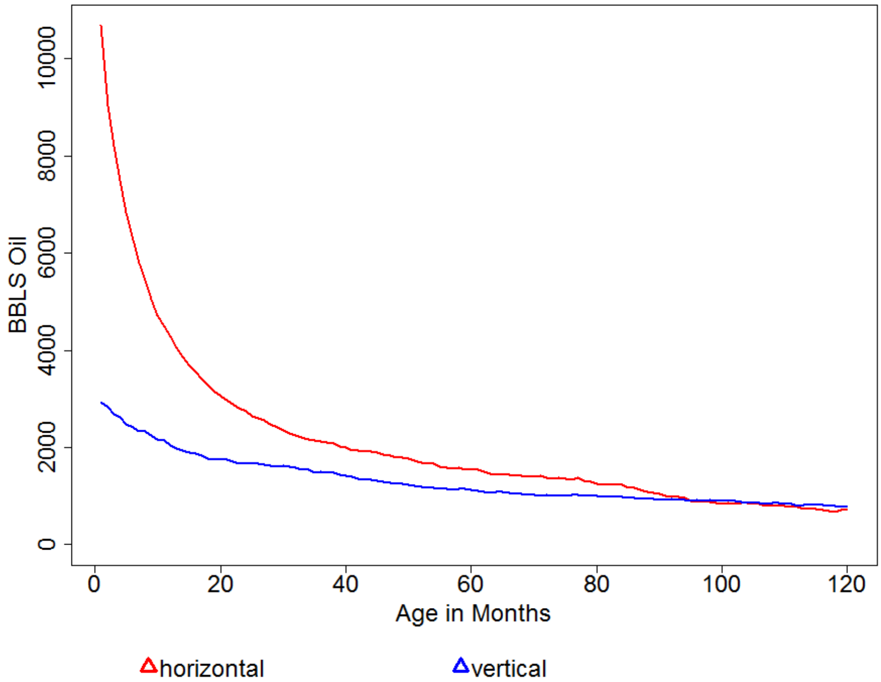

Note: Cross-sectional average of the monthly production rate of wells in the sample plotted against well age measured in months, divided into average production for horizontal and vertical wells.

Table 1. Average monthly decline rates by well type

	All wells	Conventional	Shale
All months	-0.005	-0.003	-0.016
First 12 months	-0.042	-0.001	-0.052

conventional wells. That is, the production rate of an average shale well is almost four times as high in the first month of production as that of conventional wells. Second, the average decline rate for shale wells is more than five times the average monthly decline rate for conventional wells, and within only 3-5 years the production rate is on the same level as an average conventional well. Figure A.2 in Appendix A, which shows the share of total output from shale and conventional wells over the first five years of a well's lifetime, further confirms this picture. As the figure shows, a much larger fraction of a shale well's total output is produced in the first months of the well's lifetime.

Finally, the productivity of wells, measured as the flow rate for the first month is also different across geographies. This is mainly due to geological factors and the quality of oil reservoirs. Figure A.3 in Appendix A illustrates this feature of oil extraction well. The figure displays the mean initial monthly flow rate of wells inside small geographic units in North Dakota. The darker units are the most productive.[11] Well productivity

[11] The geographic units are constructed using a 25 times 25 array raster covering all producing wells in North Dakota.

appears to be highest in the central part, in the counties of Mountrail and McKenzie. There appears to be some spatial autocorrelation of well productivity, since the darkest units rarely border the lightest units. This suggests that the main determinants of well productivity are geological characteristics, which are spatially autocorrelated at this level of disaggregation. Using a panel to analyse our question is clearly then an advantage. As we can have well fixed effects, we can control for the geographic factors as well.

2.4 Oil price indices

The oil price series used in the analysis is the West Texas Intermediate (WTI) crude oil price for delivery in Cushing, Oklahoma, provided by the Energy Information Administration (EIA). The futures prices at different time horizons are from the New York Mercantile Exchange (NYMEX), and retrieved from Datastream.

We use the log difference between the futures price for a contract with delivery at time $t + j$ and the spot price to measure the expected rate of change in the price of oil going j months ahead. Equivalently this is referred to as the futures curve slope for different time horizons. If this measure is negative, we assume that the market expects the price of oil to be lower j months from now than today's spot price. If it is positive, the price is expected to increase over j months.

Using futures prices to measure a firm's price expectations is done for several reasons. First, NYMEX futures are traded liquidly at the time horizons considered here, and with many risk-neutral traders, the futures price should equal the expected future spot price. Also, oil well operators are believed to use the futures market to make price projections. The futures prices included in this paper are real prices, so that the annual expected price change reflects real rather than nominal changes. In addition, the future prices on the NYMEX marketplace are the reference prices that physical crude oil futures market can actually be sold at in the physical market for crude oil for future delivery.

3 Empirical model and results

Our aim is to analyse to what extent oil producers respond to price incentives by changing the well completion rate and oil production (the flow) from wells. To do so we first analyse if firms adjust the flow rate from existing wells in response to price incentives (intensive margin), distinguishing between conventional and shale oil wells. Then we ask if firms optimise the completion of new wells in response to price incentives (extensive margin), comparing again the two types of well technology.

In all the analysis we make the assumption that the oil price is exogenous to the monthly output of crude oil wells and number of wells put into production at time t in North Dakota. This is plausible for several reasons. First, since there are about 500 firms in our sample, no single firm is able to exert any market power in the global market for crude oil. Second, North Dakota is not a big player in the global oil market, and the combined production of all firms is just over one million barrels of crude oil per day, which is below one per cent of the global daily output of oil. Third, the price of oil is largely determined in the global market, and the additional shale oil production coming from the US in recent years is small compared to the size of the global oil market.

3.1 Short term supply elasticity

We regress the monthly percentage change in barrels produced per well in North Dakota on two different price indices: the percentage changes in the spot price of oil, and the percentage changes in the first month spot futures spread. While the first price is the traditional indicator used for measuring supply elasticity, we use the latter price indicator as a proxy for changes in the short-term expected price trend.

For motivation, we first construct elasticity estimates using aggregate specifications. That is, rather than using the whole panel, we aggregate the output produced from the wells into two separate monthly time series; one for conventional oil production and one for shale oil production. We then run three separate regressions: A regression for conventional production, a regression for shale production and a regression for total (shale and conventional) production. Table A.2 in Appendix A displays the results. All coefficients are insignificant at the 5% level, echoing previous results in the literature that oil producers do not respond significantly to the current and near by price signals. Recall, however, that what we have done is to aggregate over individual wells. This is equivalent to imposing identical parameter values for all producing wells on either shale or conventional (or both) regardless of well or firm characteristics. This will lead to a significant loss of information about micro relations in our data set and lead to a potential aggregation bias. Only by constructing a rich panel data set, can we eliminate any potential aggregation bias over well production rates when estimating the empirical model. We turn to this now.

To this end, we estimate the short-run price elasticity of supply using the following baseline panel data specification:

$$\Delta Q_{it} = \alpha \Delta Q_{it-1} + \beta_1 \Delta P_t + \beta_2 \Delta[F_{t+1} - P_t] + \gamma age_{it} + \mu_i + \lambda_t + v_{it} \tag{1}$$

where ΔQ_{it} is the percentage change in the quantity of oil produced on a monthly frequency for well i at time t, ΔP_t is the percentage change in the WTI spot price, $\Delta[F_{t+1} - P_t]$ is the change in the spread between the first month future contract price and the current period spot price. The lagged dependent variable, ΔQ_{it-1}, is added to allow for well production changes to be autocorrelated. age_{it} is the number of months the well has been in operation, and λ_t is included to remove aggregate yearly changes in well output. Since the oil price is a macro-variable, common to all firms at time t, it is not feasible to remove all monthly fixed effects and still identify the effect of a monthly price change on output. For this reason, the model includes year-fixed effects, while μ_i is included to remove any well fixed effects. Since the model is in first differences, any well-specific characteristics such as the quality of the well bore and geological characteristics are already removed. The first-differencing also removes any time-invariant effects stemming from the drilling contractor firm or operator firm. Hence, adding μ_i to the model removes all well-specific *linear trends* over the sample period. This could be any well-specific trends in production, or well-specific technological learning. v_{it} is the idiosyncratic error term, which is clustered within the month. Since the oil price is a common shock potentially affecting all wells at time t, we allow standard errors across wells within the same time period to be correlated. In our setting, the dynamic panel bias which may be present for panels with small t is less of a concern, since t exceeds 300 time periods in our estimation.

Table 2. Supply elasticity by well type

	(Pooled regression)	(By well techn.)	(By well techn.)
Dep. var: Δ Log Prod$_{it}$			
$\Delta[P_t] * Total$	0.055*		
	(2.08)		
$\Delta[F_{t+1} - P_t] * Total$	0.082		
	(1.87)		
$\Delta[P_t]$		0.035	0.035
		(1.12)	(1.01)
$\Delta[P_t] * Shale$		0.072	0.041
		(1.28)	(0.92)
$\Delta[F_{t+1} - P_t]$		0.013	0.00
		(0.38)	(0.26)
$\Delta[F_{t+1} - P_t] * Shale$		0.220**	0.195**
		(2.87)	(2.86)
$WellFE$	No	No	Yes
$Year\,FE$	No	No	Yes
$No\,of\,Wells$	15,259	15,155	15,155
N	1,202,787	1,188,227	1,188,227
$Time\,period$	1986-2015	1986-2015	1986-2015

Note: Key results are given here. See Table A.3 in Appendix A for more details. t-statistics are in parentheses. * Significantly different from zero at the 90% level, ** 95% level and *** 99% level.

We first estimate the parameters of the model independent of extraction technology, and then identify the shale wells by an indicator variable, so as to estimate the technology-dependent parameters. The indicator variable will take the value one if it is shale wells, or zero otherwise. Given our setup, the estimates of β_1 and β_2 will capture the average response in monthly well output at time t to changes in the price level and to changes in the expected price trend, respectively. Results from the estimation of equation 1 are presented in Table **??**. Parameters are estimated from a sample of 15,259 producing wells from July 1986 to November 2015. We report here on the main parameters of interest. All parameters are displayed in Table A.3 in Appendix A.

We report the pooled regression estimates of all wells (i.e., independent of well technology) in Column 1. From this we see that the point estimate of the response for all wells to a percentage change in the spot price of oil at time t is positive at .055, nearly identical to the short term supply response found in Anderson et al. (2014) for aggregate oil production in Texas. It is, however, small in magnitude, only significant at the 10 percent level, and thus echoes the results in the literature that short-term supply elasticity for conventional oil wells is near zero. The output response to a shift in the short-term expected price trend is also positive (0.082) but not significant. Still. the magnitude of the estimate suggests that producers respond to forward-looking price indices slightly more so than to shifts in the current price level.

Moving to Column 2 and 3, we now add the indicator variable that takes the value of one for shale wells and zero for conventional wells in order to estimate separate parameters according to technology. Column 3 differs from column 2 in that we also add well fixed effects and year fixed effects to control for well-specific trends and year aggregate effects. As can be seen, using either specification in column 2 or 3, now the responses of conventional and shale wells differ significantly. In particular, we find that the responses of conventional wells to both changes in the price level and to changes in the spot future spread are small and statistically insignificant. Turning to shale wells, however, the response to current price changes is somewhat larger. From column 3 we find the response to be 0.76, just statistically significant at the 5 percent level.[12] Even larger is the output response of shale wells to a shift in the one-month spot future spread. The immediate supply response to a ten percent increase in the spot future spread, or short-term expected price trend, is to increase the monthly well flow by close to two percent. The response is significant at the 1 percent level.[13] For an average shale well that has produced for less than a year, this means increasing the monthly flow rate by nearly 110 barrels per day following a one standard deviation increase in the spread.

So far we have investigated supply elasticity with respect to spot prices and short-term price changes only. The futures market for crude oil provides intertemporal price signals, that producers also consider closely. In particular, producers may reduce/increase output today when they expect prices a year from now to be higher/lower than today's spot price. To investigate more formally whether firms respond differently to shifts to long-term price expectations, we add long-term future spreads to the baseline model. If the spread shifts

[12]We can reject the null that the absolute elasticity for share wells (0.035 + 0.041) to changes in price level is zero at the 5 percent confidence level; (Prob > F = 0.0467).

[13]We can reject the null that the absolute elasticity for share wells (0.00 + 0195) to changes in the sport future spread is zero at the 1 percent confidence level; (Prob> F = 0.0052).

Table 3. Supply elasticity with forward-looking prices

	(3-month horizon)	(6-month horizon)	(12-month horizon)
Dep. var: Δ Log Prod_{it}			
$\Delta[P_t]$	0.051	0.036	0.032
	(1.69)	(1.02)	(0.82)
$\Delta[P_t] * Shale$	-0.054	-0.054	-0.051
	(-1.20)	(-1.03)	(-0.88)
$\Delta[F_{t+1} - P_t]$	0.013	0.007	0.006
	(0.35)	(0.21)	(0.14)
$\Delta[F_{t+1} - P_t] * Shale$	0.15***	0.17***	0.19***
	(3.04)	(3.30)	(3.44)
$\Delta[F_{t+3} - P_t]$	0.002		
	(0.01)		
$\Delta[F_{t+3} - P_t] * Shale$	-1.456***		
	(-3.31)		
$\Delta[F_{t+6} - P_t]$		0.042	
		(0.22)	
$\Delta[F_{t+6} - P_t] * Shale$		-0.820***	
		(-2.66)	
$\Delta[F_{t+12} - P_t]$			-0.031
			(0.25)
$\Delta[F_{t+12} - P_t] * Shale$			-0.464**
			(-2.13)
$Well FE$	Yes	Yes	Yes
$Year\ FE$	Yes	Yes	Yes
$No\, of\, Wells$	15,155	15,155	15,155
N	1,188,227	1,188,227	1,188,227
$Time\, period$	1986-2015	1986-2015	1986-2015

Note: Key results are given here. See Table A.4 in Appendix A for more details. t-statistics are in parentheses. * Significantly different from zero at the 90% level, ** 95% level and *** 99% level.

upwards, the price is expected to increase at a faster rate, while if it shifts downwards, the rate of price growth is expected to be slower. If so, we should expect to see that shale wells on average reduce output when spreads grow larger, while we expect no relationship for conventional wells. The corresponding augmented model now reads:

$$\Delta Q_{it} = \alpha \Delta Q_{it-1} + \beta_1 \Delta P_t + \beta_2 \Delta[F_{t+1} - P_t] + \beta_3 \Delta[F_{t+j} - P_t] + \gamma age_{it} + \mu_i + \lambda_t + v_{it} \quad (2)$$

where we have added a term gauging the spot future spread at different time horizons, i.e., three-, six- and twelve-months spreads. Due to the high correlation between the spreads, we add them one at the time. Key results are presented in Table 3. Throughout, we control for well and year fixed effects, as well as well age, but for brevity we omit them from this table. Detailed results are reported in Table A.4 in Appendix A.

Starting with Column 1 in Table 3, we have, in addition to the same price indices as in model 1, also added changes in the three-month spot future spread for both technologies. As can be seen, the results from Table **??** are more or less reiterated in Table 3, although the coefficient for shale wells' response to the one-month spot future spread is slightly reduced, albeit still significant.[14] For changes in the three-month expected price trend, however, the responses for the two types of wells differ greatly. While conventional wells show no significant response to a shift in the three-month spot future spread, wells using shale technology respond significantly negative to the same positive shift in the price spread.

In Column 2, we add instead the six-month expected price trend to the baseline model. Again, the estimated effect of shifts in the one-month spot future spread are found to be close to previous results. However, as can be seen, there is again a large negative response for shale wells to shifts in the longer term spot future spread, with no response from conventional wells. Hence, we believe that well operators also effectively use the ground as a storage facility at the six-month horizon when price growth is expected to be higher, or increase production when price growth is expected to be lower. Finally, in the last Column, the twelve-month spread is added to the baseline model. Again, shale well operators respond negatively to a positive shift in the twelve-month spot future spread, while no significant response is found for conventional wells.

Overall, conventional oil wells respond insignificantly to the price indices included in our models. Shale wells, however seem to respond positively to increases in short-term expected price trends, but negatively to increases in longer-term expected price trends. The results strongly indicate that shale wells are more price-elastic than conventional wells. Additionally, it appears that firms exploit the inherent flexibility of shale technology to allocate production volumes intertemporally and to engage in short-term below-ground storage.

3.2 Well completion - firm level analysis

Having analysed the response of well production to price signals, we want to examine if firms' decision to enter new wells into their stock of producing wells is also affected by price signals. The decision to complete a well is equivalent to exercising the real option to produce, since this is the actual start of production from a well. As explained above, the timing of when to *produce* is different from the timing of *drilling* the well. The reason

[14] Responses to changes in the current price level, however, is no longer significant.

is that wells are frequently drilled, but left uncompleted for a some time. Once a well is completed, however, it starts producing. Thus, we define a well as completed the moment they have their first month of production.

We postulate that shale firms more often than conventional wells leave wells that are drilled uncompleted, or in below ground storage. These are assets than can be quickly liquidated when markets are favorable. The steep production profile of shale wells (c.f., Figure 3) provides an incentive for optimizing the timing of completion. Hence, we should expect to see a delay between the time a well is spudded (drilled) and the first production month. Figure 4 illustrates this well. It displays the empirical probability function of delay between spudding a well and the first month with production, measured in number of months. The figure illustrates that on average there is a longer gap between drilling/spudding for shale wells than for conventional. In addition, the distribution is more dispersed; hence variation is larger. In sum, this supports the idea that deciding to complete a well is a choice variable, which motivates why we study supply elasticity on the extensive margin.

Figure 5 further motivates such a study. It shows a correlation plot of the twelve-month spot future spread against the monthly number of well completions (in logs) since January 2009. The left plot shows the relationship for monthly conventional well completions and the right plot for shale well completions. As can be seen, the monthly completions of shale wells are strongly negatively correlated with the twelve-month spot future spread. Conventional well completions, however, have a zero correlation and show no such relationship. The negative correlation between monthly shale well completions and the annual spot future spread implies that shale firms reduce/increase the monthly number of wells they complete when the future curve is upward/downward sloping. The vertical pattern in the scatter plot is due to the low number of conventional wells completed during this period. Therefore, in order to check if there is any relationship when including a larger number of observations, we have constructed the same plot including all conventional wells since 1983. Figure A.4 in Appendix A shows this plot, which strongly suggests that a negative relationship has never existed between monthly conventional well completions and the annual price spread.

Hence, if there is a negative relationship between firms' monthly well completions and an increase in the medium- to long-term spot future spread, it implies that firms reduce the rate of completions when the future spot price is expected to be higher than the spot price today, effectively leaving oil in the ground awaiting a higher expected price. To test this more formally, we set up a firm-level panel data model. In order to identify the potential heterogeneous effect of changes in some price indices on firms' well completion rate, depending on the use of shale technology versus conventional extraction technology, we calculate each firm's share of shale wells as a fraction of its total stock of producing wells, and denote this θ. This measure proxies how much shale acreage a firm has access to and the intensity of shale technology in its stock of wells. This measure is between

Figure 4. Drilling and completions.

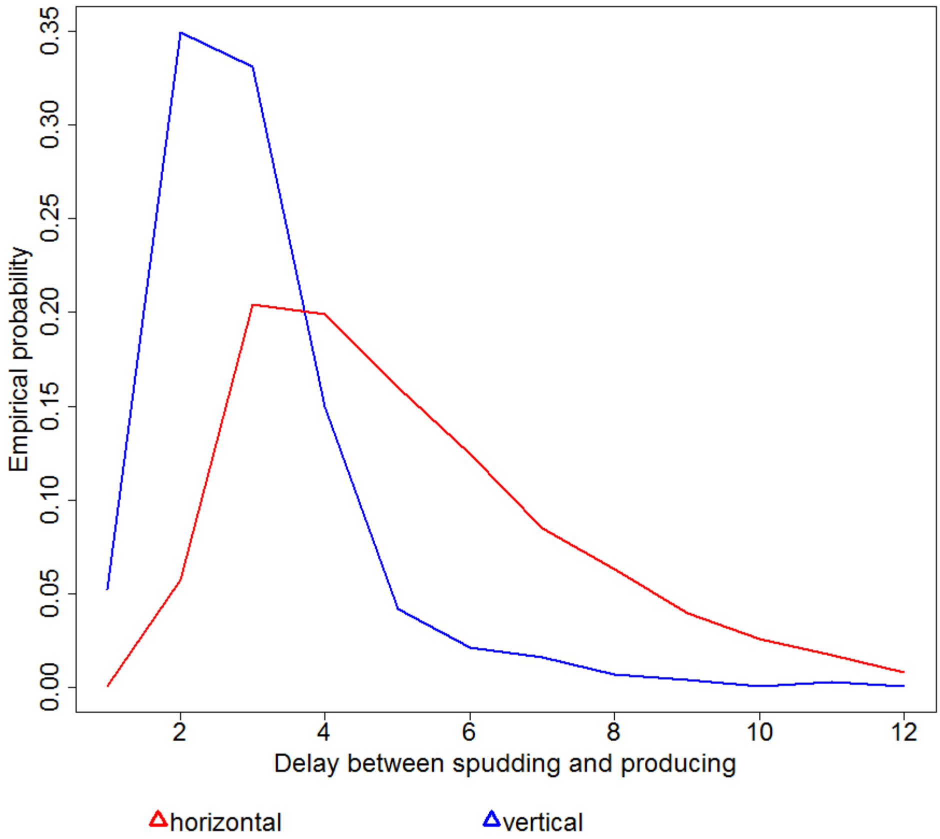

Note: This figure displays the empirical probability function of delay between spudding a well (drilling) and the first month with production, measured in number of months. The blue line indicates conventional wells while the red one focusses on unconventional. The sample consists of all wells that where spudded after June 1986, which produced oil at least one month before November 2015.

zero and one, and it varies both between firms and within firms over time. Figure A.5 in Appendix A details the time-varying share of shale wells for a sample of 20 firms in our data. As can be seen, many of the firms used both shale and conventional technology during the sample period. Still, a majority of the 20 firms depicted transitioned quite heavily towards shale technology.

Since the regression model predicts well completions in first-differences, any stable confounding firm characteristics or firm fixed effects are removed. This addresses the issue that, for example, different types of firms could be using different technology. It also removes firms with zero change in extraction technology during the sample period from the estimation. Other stable confounders include firm ownership structure, average size, location of headquarters or access to capital markets. Since we nevertheless include firm fixed effects in the estimation, we also remove any firm-specific trends in well completions, such as particular firms having adapted the new technology at a faster rate than others, or firm-specific trends in cost levels, revenues or any other non-observable trends. The

Figure 5. Monthly well completions and expected annual oil price change

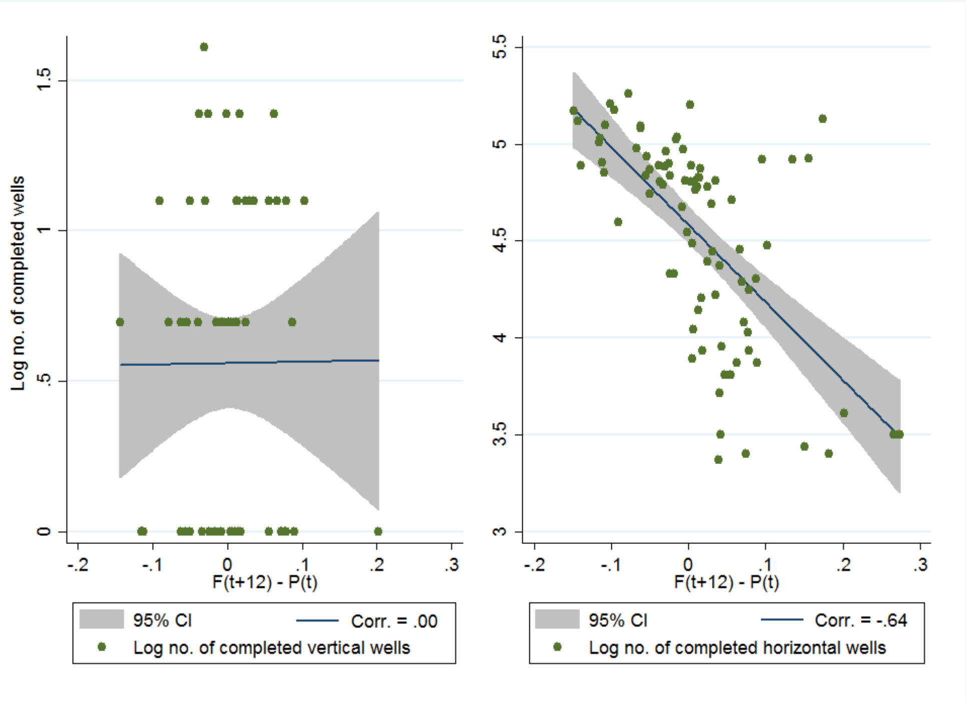

Note: The Figure shows the correlation between expected annual change in the price of crude oil (spread between the twelve-month future contract and the spot contract) and the mean monthly well flow of conventional oil wells (left) and shale oil wells (right). The time period is 1.2009-11.2015.

between-firm variation of firms' shale share allows us to add monthly fixed effects and thus remove the common trend in firms' use of shale technology in North Dakota since the technology arrived. Combined, this enables us to identify whether firms using more shale technology than the average firm at time t change their well completion rate in response to the same price shocks differently from firms with less exposure to shale technology.

To define the slope coefficient of interest, we do as follows. Let $\Delta\dot{W}_{it}$ be the within-firm change in the monthly number of completed wells by firm i, demeaned to remove any unobserved firm-level trends. From this, we also remove the monthly fixed effects, i.e., the average change in well completions across firms at time t; denoted $\Delta\bar{W}_t$. Consequently, $\Delta\ddot{W}_{it}$ is the (remaining) variation in firm-level well completion rates:

$$\Delta\ddot{W}_{it} = \Delta\dot{W}_{it} - \Delta\bar{W}_t$$

Furthermore, let $\dot{\theta}_{it-1}$ be the demeaned within-firm variation in the fraction of wells that are shale wells at time *t-1* and $\bar{\theta}_{t-1}$ the average share of shale wells across all firms at time *t-1*. Figure A.5 in Appendix A details the evolution of θ_i for a selected group of companies over the sample time horizon.[15] As in the case of well completions, removing the monthly fixed effects amounts to subtracting the average firm's shale share in period

[15] We use the share at time *t-1*, since otherwise the measure will contain completed wells in period t.

t-1, so that $\ddot{\theta}_{it-1} = \dot{\theta}_{it-1} - \bar{\theta}_{it-1}$. Assume P to be some measure of price. The slope that we are interested in estimating is then

$$\beta = \frac{\partial(\Delta \ddot{W}_{it})}{\partial(\ddot{\theta}_{it-1})P}$$

The coefficient captures the marginal response of firm i's well completions (relative to the average firm's well completions) to a price change, depending on the share of shale wells used in extraction by firm i (relative to the average firm's shale well share). Assuming the price index used is the spot future term spread, a negative slope would indicate that when a firm has a higher than average share of shale wells in its stock of producing wells, a positive price spread induces the firm to reduce well completions relative to the average firm's completions. The full empirical specification will then be

$$\Delta\ddot{W}_{it} = \alpha\Delta\ddot{W}_{it-1} + \beta_1\ddot{\theta}_{it-1}[F_{t+1} - P_t] + \beta_2\ddot{\theta}_{it-1}[F_{t+j} - P_t] + \gamma stock_{it-1} + \epsilon_{it} \qquad (3)$$

where $\Delta\ddot{W}_{it}$ is the within-firm percentage change in the monthly number of completed wells in North Dakota, relative to the average firm's completions at time t. $F_{t+1} - P_t$ is the one-month spot future spread, and $F_{t+j} - P_t$ is the spot futures spread j months ahead. In the estimations, we will use the rate of expected price change three-, six-, twelve- and 18- periods ahead. $stock_{it-1}$ is the total stock of wells operated by a firm in period *t-1*, and ϵ_{it} is the error term, which is allowed to be clustered within the month.

Table 4, Column 1 shows the results for the three-month spot future spread when utilising the full sample of 366 firms for the time period from 1986 to 2015. The effect of a one standard deviation positive shock in the three-month spot future spread induces a firm to *reduce* its monthly well completion rate immediately by almost five percent. Hence, firms appear to respond significantly to the slope of the three-month future curve, and more so the more intensively they use shale technology.

In Columns 2, 3 and 4 we explore the effects across longer term structures. The response from firms' monthly well completions are significant and negatively associated with the price spread for each of the time horizons, confirming the initial negative relationship depicted by the simple correlation in Figure 5. By symmetry, the sign of the point estimates indicate that when the term spreads are negative and the oil market is in backwardation, firms respond by immediately *increasing* their monthly well completions by the same magnitudes. These estimates indicate that firms respond to the term structure of oil prices by adjusting the rate of completion. The size of the existing stock of wells in the previous period for a firm predicts weakly, but not economically significantly, the change in current period well completions. Throughout the regression results, the short-term price spread is insignificant as explanatory variable for monthly well completions. This indicates that producers are mostly forward-looking and are less concerned about the price spread variation in the very short term.

To sum up, we have documented that shale wells reduce/increase the number of new wells (extensive margin) they add to the stock of producing wells substantially today

Table 4. Firm panel: monthly well completions

	(3 mo. hor.)	(6 mo. hor.)	(12 mo. hor.)	(18 mo. hor.)
Dep. var: ΔWells_{it}				
$[F_{t+1} - P_t] * ShareShale$	0.0636	0.0736	0.0784	0.0844
	(0.94)	(1.10)	(1.17)	(1.23)
$[F_{t+3} - P_t] * ShareShale$	-0.476**			
	(-2.78)			
$[F_{t+6} - P_t] * ShareShale$		-0.281**		
		(-3.28)		
$[F_{t+12} - P_t] * ShareShale$			-0.173**	
			(-3.05)	
$[F_{t+18} - P_t] * ShareShale$				-0.136**
				(-2.91)
Time FE	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
N	57187	57187	57187	57187
Firms	366	366	366	366
Time period	1986-2015	1986-2015	1986-2015	1986-2015

Note: Key results are given here. See Table A.5 in Appendix A for more details. t-statistics are in parentheses. * Significantly different from zero at the 90% level, ** 95% level and *** 99% level. $\Delta[F_{t+j} - P_t] * ShareShale$ is the percentage change in the difference between the future price at time horizon j and the spot price at time t, interacted with the share of shale wells in the total stock of wells.

when the long-term (three-, six-, twelve- or 18-months) future curve is upward sloping/downward sloping. The numbers are not trivial. For a ten percent increase in the twelve-month spot future spread, producers respond by reducing the number of wells they complete by close to 2 percent, or an average of eight wells pr month. This is equivalent to holding back about 80,000 barrels of daily crude oil production from the market.[16] We find no evidence of such dynamics for conventional oil producers.

[16] This elasticity is based on an unconditional least squares estimation, and a first month average production of 10,000 barrels.

4 Robustness

We analyse robustness along several dimensions. Along the intensive margin, we do two extensions. First, as documented many times above, conventional and shale oil wells differ markedly in their production persistence. This suggests that we could also allow ϕ (the coefficient on lagged production) to differ across the two types of wells in equation (1). Doing so, we find ϕ to be significantly different across the two types of wells. Still, the main results are robust and can be obtained at request. Second, we estimate separate regressions for conventional and shale wells over different sample periods (1986-2015 for conventional wells and 1990-2015 for unconventional wells). Results can be seen in Table A.6 in Appendix A. Again, we see that only shale wells respond significantly positive to short term future spread (at the 5 percent level).

For the firm level analysis (the extensive margin), all results can be viewed in Table A.7 in Appendix A. In all the estimations, we regress the number of wells completed on the one-month and the six-month price spreads using equation (3), but doing several alternative specifications. First, as discussed above, the number of wells completed across firms in the sample period is highly concentrated, with 60 percent of firms having completed only one to three wells during the sample period. Since a large portion of the firms in North Dakota are small units operating only one or two wells, a substantial share of the firms included in the main regression show little variation in the monthly number of wells completed. Therefore, in order to achieve a quantitatively more correct estimate of the effect of the price spread on firms' completions, we exclude firms that have completed two or fewer wells during the sample period. Many of these firms are non-professional, family-owned firms, adding up to a total of 145 firms. Results can be viewed in Column 1 in Table A.7, where we regress the number of wells completed on the six-month and one-month price spread. As can be seen, excluding this portion of firms increases the marginal effect of a unit increase in the six-month price spread by about 50 percent compared with the baseline regression. For a ten percent increase in the six-month price spread, firms with only shale wells reduce the number of wells they complete by an average of 3-4 percent relative to the average monthly change in the number of wells completed by all firms. Also for the one-month price spread there is a significant, positive response by firms, indicating that they increase the completions when the short-term spread increases.

Another concern is that the estimates we obtain for shale firms are not due to the intensity of shale technology used *within* the firm, but rather the average difference in responses between highly specialised shale firms and specialised conventional firms. In order to address this, we exclude firms in periods when they have a shale well share equal to one, or at the time they became specialised shale firms. Column 2 displays the results. The coefficient for the six-month price spread increases in magnitude relative to the baseline, which strongly indicates that the effect of a price spread on firms' monthly well completions is strongest for the firms that have both conventional and shale well technology within the firm. Also, the magnitude of the response to the one-month price spread is larger.

A third concern is that there may be regulatory shocks in certain time periods that only affect shale wells. For this to bias estimates, the shale-specific shocks would also have to be correlated with oil price changes. An example of such a shock could be regulation concerning the amount of water and chemicals allowed for use in the fracking operations. Such shale well specific time fixed effects would, if not controlled for, end up in the error term as an omitted variable. In order to address this concern, we add time fixed effects interacted with the average share of shale wells across firms, but within the same time period. This will allow the aggregate time fixed effects to have a differential impact, depending on the share of shale wells across firms in the time period. As can be seen in Column 3, the coefficient on the six-month price spread is still significant and large. The one-month price spread is also large and significant when controlling for the shale specific time fixed effects.

Finally, we want to obtain estimates of the effect of the price spread on firms' completions during the phase when shale technology was widely adapted and firms completed many wells. After all, it is in this period that shale technology was most widespread in North Dakota. Therefore, in the last column of Table A.7, we restrict the sample period analysed to start in January 2009. At the same time, we exclude the 20 percent of firms that only completed one or maximum two wells during this period. An advantage of obtaining estimates for this period is that we remove the one time effect of the large oil price shock in 2008 on our coefficients. As can be seen, it is in the time period where shale technology was the most widely used and for firms that were the most active, that the effect is the strongest. For a 10 percent increase in the six-month price spread, firms on average reduce their completions by more than 11 percent in this time period. By symmetry, the same firms increase the completion of wells if the price spread declines by 10 percent.

5 Interpretation - The Hotelling rule

Are the reported results for shale oil producers in line with Hotelling (1931)'s model of optimal exhaustible resource extraction? Reserves are an inventory, and as such the decision to produce is an intertemporal choice of when to convert below-ground inventory to above-ground inventory, stored in pipelines, oil terminals and refinery facilities.

The Hotelling (1931) model assumes resource owners are forward-looking and maximise wealth by trading extraction today for extraction in the future. An implication of the model is that in a simple equilibrium with no production cost or risk, expected price growth for exhaustible commodities such as crude oil matches the opportunity cost of storage - that is, the nominal interest rate (r). Crucially, Hotelling's price path is only achieved when producers enjoy complete flexibility regarding when to produce. This

Figure 6. Expected *net* price growth and shale well output

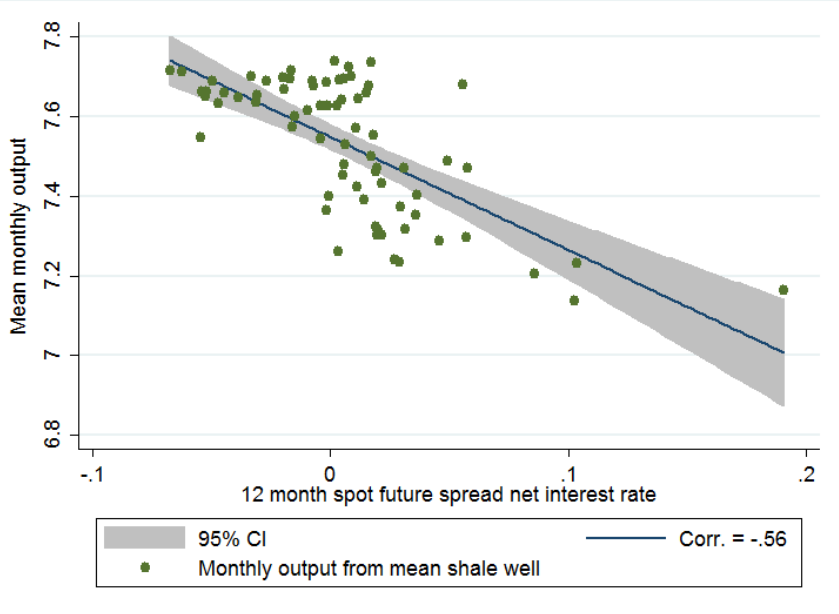

Note: Twelve-month spot future spread net of annual risk-free interest rate on x-axis, mean shale well output (log) on y-axis.

is hard to reconcile with *conventional* oil production (see for instance Anderson et al. (2014)). Our empirical results, however, indicate that the degree of output flexibility depends on the technology applied and that firms using shale oil technology are more flexible in allocating output intertemporally. For producers to fully comply with the classic theory, one should then observe that production is reduced to zero when expected price growth exceeds the nominal interest rate, and vice versa when prices are expected to grow below the risk-free rate. Although well flow rates in our data are rarely reduced to zero, using the futures price as a proxy for the expected spot price, our results indeed suggest that shale well technology allows producers to allocate production more in line with what the classic Hotelling theory predicts as the optimal response to deviations from the Hotelling price path.

To relate more directly to Hotelling's optimal extraction rule, we construct price indices net of interest rates. Such a measure should capture short-term expected deviations from the Hotelling price path. To achieve this, we define expected annual price growth net of interest rates as the twelve-month spot future spread less the three-month US treasury rate, annualised to reflect the expected return from putting one dollar into bonds for a year, as opposed to deferring output equal to one dollar and earning an expected return equal to expected annual price growth. Still, as the risk-free interest rate since 2008 has remained close to zero, we expect the results to show similar patterns as in the main analysis reported in Section 3.

And they do. As can be seen from Figure 6, when expected annual price growth exceeds the nominal interest rate, shale oil wells in North Dakota on average have a reduced flow rate. However, when expected annual price growth is below the nominal interest rate or negative, shale wells on average have a higher flow rate. The pattern shown in Figure 6 is consistent with Hotelling's theory of how resource extraction firms should respond to deviations from Hotelling's price path. To show this more formally, we also construct *net* oil price spreads for the three- and six-months horizons (in addition to the twelve-month horizon),[17] and reiterate the results from Table A.4. The results are shown in Table A.8 in Appendix A. As expected, the results are very similar to the baseline results, confirming that results for shale firms can be interpreted in line with Hotelling's optimal extraction rule.

6 Conclusion

This paper has shown that output flexibility in oil production depends on the extraction technology. Furthermore, we have quantified large and substantial differences in the response of crude oil output to price signals between conventional and shale oil wells in North Dakota. In particular, shale wells increase the monthly production in shale wells by nearly two percent for a ten percent increase in the spot future spread. Furthermore, when the price one year ahead is expected to be higher than today's price, operators of shale wells reduce both the number of wells they put into production and the output rate from producing wells. The magnitudes are of economic importance. For a ten percent increase in the twelve-month spot future spread, producers respond by reducing the number of wells they complete by an average of eight wells. This is equivalent to holding back about 80,000 barrels of monthly crude oil production from the market. We find no such evidence for conventional oil. Finally, we document that the output pattern of shale oil wells is consistent with the Hotelling theory of optimal extraction. For the behaviour of conventional oil wells, we find no such evidence.

This study is the first to investigate the price elasticity of tight oil production and to compare it with conventional oil production. As the share of unconventional oil production increases, it is important to have good estimates of price elasticity of supply. It is also necessary to have knowledge about the production behaviour of shale oil in general, as it may affect a wide range of outcomes such as local labour market dynamics, tax policies, oil supplies and prices.

[17] As a proxy for the interest rate, we use the effective federal funds rate, compounded to match the twelve-month and six-month spreads.

References

Anderson, S. T., R. Kellogg, and S. W. Salant (2014). Hotelling under pressure. *Journal of Political Economy* (Forthcoming).

Baumeister, C. and J. D. Hamilton (2015). Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks. Manuscript, University of Notre Dame and UCSD.

Caldara, D., M. Cavallo, and M. M. Iacoviello (2016). Oil price elasticities and oil price fluctuations. International Finance Discussion Papers 1173.

Carpenter, C. (2014). Design optimization of horizontal wells with multiple hydraulic fractures in the bakken shale. *Journal of Petroleum Technology 66*(11), 118–123.

Clerici, A. and G. Alimonti (2015). Oil shale, shale oil, shale gas and non-conventional hydrocarbons. In *EPJ Web of Conferences*, Volume 98. EDP Sciences.

Dahl, C. and M. Yücel (1991). Testing alternative hypotheses of oil producer behavior. *The Energy Journal 12*(4), 117–138.

Devold, H. (2013). *Oil and gas production handbook: an introduction to oil and gas production, transport, refining and petrochemical industry*. ABB AS.

Dunn, S. (2016). Fracking 101: Breaking down the most important part of today's oil, gas drilling. *The Tribune, Greely and Weld County, Colorado (October 14, 2016)* (URL http://www.greeleytribune.com/news/local/fracking-101-breaking-down-the-most-important-part-of-todays-oil-gas-drilling).

Gaswirth, S. B., K. R. Marra, T. A. Cook, R. R. Charpentier, D. L. Gautier, D. K. Higley, T. R. Klett, M. D. Lewan, P. G. Lillis, and C. J. Schenk (2013). Assessment of undiscovered oil resources in the Bakken and Three Forks Formations, Williston Basin Province, Montana, North Dakota, and South Dakota, 2013. Technical report, US Geological Survey.

Gopinath, S. (2014). Factbox: Breakeven oil prices for U.S. shale: Analyst estimates. Reuters. retrieved from http://www.reuters.com.

Griffin, J. M. and D. J. Teece (2016). *OPEC behaviour and world oil prices*. Routledge.

Hamilton, J. D. (2009). Causes and consequences of the oil shock of 2007–08. *Brookings Papers on Economic Activity* (1), 215–283.

Hogan, W. W. (1989). *World oil price projections: a sensitivity analysis*. Harvard University, Energy and Environmental Policy Center, John F. Kennedy School of Government.

Hotelling, H. (1931). The economics of exhaustible resources. *Journal of Political Economy 39*(2), 137–175.

Hubbert, M. K. et al. (1956). *Darcy's law and the field equations of the flow of underground fluids*. Shell Development Company, Exploration and Production Research Division.

Kilian, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *The American Economic Review 99*(3), 1053–1069.

Kilian, L. and D. P. Murphy (2012). Why agnostic sign restrictions are not enough: understanding the dynamics of oil market var models. *Journal of the European Economic Association 10*(5), 1166–1188.

Lippi, F. and A. Nobili (2012). Oil and the macroeconomy: a quantitative structural analysis. *Journal of the European Economic Association 10*(5), 1059–1083.

Meissner, F. F. (1984). Petroleum geology of the Bakken Formation Williston Basin, North Dakota and Montana.

Miller, B. A., J. M. Paneitz, M. J. Mullen, R. Meijs, K. M. Tunstall, and M. Garcia (2008). The successful application of a compartmental completion technique used to isolate multiple hydraulic-fracture treatments in horizontal Bakken shale wells in North Dakota. In *SPE Annual Technical Conference and Exhibition*. Society of Petroleum Engineers.

North Dakota Studies (2016). Module 2: Petroleum & Natural Gas: A history of fossil fuels in North Dakota. Energy: Powered by North Dakota,, Official Portal for North Dakota State Government. Retrieved from: http://ndstudies.gov/energy/level1/module-2-petroleum-natural-gas/history-fossil-fuels-north-dakota.

Pesaran, M. H. (1990). An econometric analysis of exploration and extraction of oil in the UK continental shelf. *The Economic Journal 100*(401), 367–390.

Ramcharran, H. (2002). Oil production responses to price changes: an empirical application of the competitive model to OPEC and non-OPEC countries. *Energy Economics 24*(2), 97–106.

Smith, J. L. (2009). World oil: market or mayhem? *The Journal of Economic Perspectives 23*(3), 145–164.

Appendices

Appendix A Tables and Figures

Table A.1. Summary Statistics: This table provides summary statistics for the time series of monthly data, the cross-section of wells, and individual observations on well-month level. Mean well production is in barrels per month

Time Series		**St.dev**	**Min**	**Max**
Start	1986-07-01			
End	2015-11-01			
No. of months	352			
Mean well production	1747.1	2141.3	0	45201.5
Mean log change well production	-0.019	0.579	-5.46	5.93
Mean well completions per firm	0.2	1.2	0	32.1
Mean spot price	54.25	25.03	11.15	126.22
Mean price change	-0.0059	0.09	-0.41	0.44
Mean Twelve-month spread	-0.002	0.06	-0.19	0.228
Total cross-section				
No. of wells	15,485			
No. of pools	34			
No. of firms	489			
No. of fields	590			
Observations	1,256,191			
Shale wells		**St.dev**	**Min**	**Max**
No. of wells	11,735			
Mean well production (bls)	3350.7	3006.9	0	45201
Mean change well production	-0.029	0.69	-9.12	9.59
Mean well age (months)	42	32	0	163
Conventional wells		**St.dev**	**Min**	**Max**
No. of wells	3,750			
Mean well production (bls)	684.1	1330.9	0	23028
Mean change well production	-0.013	0.49	-7.43	9.4
Mean well age (months)	252	88.5	0	682

Table A.2. Aggregate supply elasticity by production technology

	(Pooled)	(Conventional)	(Shale)
Dep. var: Δ Log Prod_t			
Δ Log Prod_{t-1}	-0.457***	-0.558***	-0.237***
	(-11.13)	(-15.34)	(-3.62)
$\Delta[P_t]$	0.038	0.043*	0.03
	(1.15)	(1.81)	(0.69)
$\Delta[F_{t+1} - P_t]$	0.0153	-0.00104	0.006
	(0.45)	(-0.03)	(0.13)
Constant	0.009	-0.007	0.022
	(3.88)	(-3.42)	(6.23)
N	358	358	311
Time period	1986-2015	1986-2015	1990-2015

Note: Regressions are estimated using Newey-West standard errors robust to heteroscedasticity and first order autocorrelation, with one lag included. t-statistics are in parentheses. * Significantly different from zero at the 90% level, ** 95% level and *** 99% level.

Table A.3. Supply elasticity by well type

	(Pooled regression)	(By well technology)	(By well technology)
Dep. var: Δ Log Prod$_{it}$			
Δ Log Prod$_{t-1}$	-0.3350***	-0.335***	-0.338***
	(-74.84)	(-74.62)	(-71.61)
Well age$_{it}$	0.000***	0.000***	0.001
	(11.78)	(12.36)	(1.20)
$\Delta[P_t] * Total$	0.055*		
	(2.08)		
$\Delta[F_{t+1} - P_t] * Total$	0.082		
	(1.87)		
$\Delta[P_t]$		0.035	0.035
		(1.12)	(1.01)
$\Delta[P_t] * Shale$		0.072	0.041
		(1.28)	(0.92)
$\Delta[F_{t+1} - P_t]$		0.013	0.00
		(0.38)	(0.26)
$\Delta[F_{t+1} - P_t] * Shale$		0.220**	0.195**
		(2.87)	(2.86)
$Well FE$	No	No	Yes
$Year\,FE$	No	No	Yes
$No\,of\,Wells$	15,259	15,155	15,155
N	1,202,787	1,188,227	1,188,227
$Time\,period$	1986-2015	1986-2015	1986-2015

Note: t-statistics are in parenthesis. * Significantly different from zero at the 90% level, * 95% level, *** 99% level. The dependent variable is the monthly percentage change in the production by well i. Well age$_{it}$ is the number of months since the first production month for well i at time t. $\Delta[P_t]$ is the percentage change in the spot price of oil, while the spot future spread $\Delta[F_{t+j} - P_t]$ is the percentage change in the difference between the future price at time horizon j and the spot price at time t. θ_i is an indicator variable equal to one for wells using shale technology.

Table A.4. Supply elasticity with forward-looking prices

	(3-month horizon)	(6-month horizon)	(12-month horizon)
Dep. var: Δ Log Prod_{it}			
Δ Log Prod_{t-1}	-0.338***	-0.337***	-0.337***
	(-71.63)	(-73.62)	(-73.65)
Well age_{it}	0.001	0.000	0.00
	(1.39)	(4.42)	(4.41)
$\Delta[P_t]$	0.051	0.036	0.032
	(1.69)	(1.02)	(0.82)
$\Delta[P_t] * Shale$	-0.054	-0.054	-0.051
	(-1.20)	(-1.03)	(-0.88)
$\Delta[F_{t+1} - P_t]$	0.013	0.007	0.006
	(0.35)	(0.21)	(0.14)
$\Delta[F_{t+1} - P_t] * Shale$	0.15***	0.17***	0.19***
	(3.04)	(3.30)	(3.44)
$\Delta[F_{t+3} - P_t]$	.002		
	(0.01)		
$\Delta[F_{t+3} - P_t] * Shale$	-1.456***		
	(-3.31)		
$\Delta[F_{t+6} - P_t]$		.042	
		(0.22)	
$\Delta[F_{t+6} - P_t] * Shale$		-0.82***	
		(-2.66)	
$\Delta[F_{t+12} - P_t]$			-0.031
			(0.25)
$\Delta[F_{t+12} - P_t] * Shale$			-0.464**
			(-2.13)
$Well FE$	Yes	Yes	Yes
$Year\, FE$	Yes	Yes	Yes
$No\, of\, Wells$	15,259	15,155	15,155
N	1,202,787	1,188,227	1,188,227
$Time\, period$	1986-2015	1986-2015	1986-2015

Note: t-statistics are in parenthesis. * Significantly different from zero at the 90% level, * 95% level, *** 99% level. The dependent variable is the monthly percentage change in the production by well i. Well age_{it} is the number of months since the first production month for well i at time t. $\Delta[P_t]$ is the percentage change in the spot price of oil, while the spot future spreads $\Delta[F_{t+j} - P_t]$ is the percentage change in the difference between the future price at time horizon j and the spot price at time t. θ_i is an indicator variable equal to one for wells using shale technology.

Table A.5. Firm panel: monthly well completions

	(3-mo. hor.)	(6-mo. hor.)	(12-mo. hor.)	(18-mo. hor.)
Dep. var: Δ Wells$_{it}$				
Δ Wells$_{it-1}$	-0.472***	-0.472***	-0.472***	-0.471***
	(-23.04)	(-23.03)	(-23.04)	(-22.55)
Stock of wells$_{t-1}$	-0.00582	-0.00586	-0.00594*	-0.00702*
	(-1.92)	(-1.94)	(-1.97)	(-2.07)
$[F_{t+1} - P_t] * ShareShale$	0.0636	0.0736	0.0784	0.0844
	(0.94)	(1.10)	(1.17)	(1.23)
$[F_{t+3} - P_t] * ShareShale$	-0.476**			
	(-2.78)			
$[F_{t+6} - P_t] * ShareShale$		-0.281**		
		(-3.28)		
$[F_{t+12} - P_t] * ShareShale$			-0.173**	
			(-3.05)	
$[F_{t+18} - P_t] * ShareShale$				-0.136**
				(-2.91)
$Time\,FE$	Yes	Yes	Yes	Yes
$Firm\,FE$	Yes	Yes	Yes	Yes
N	57187	57187	57187	57187
$Firms$	366	366	366	366
$Time\,period$	1986-2015	1986-2015	1986-2015	1986-2015

Note: t-statistics are in parenthesis. * Significantly different from zero at the 90% level, * 95% level, *** 99% level. The dependent variable is the monthly percentage change in the number of wells completed by firm i. Stock of wells$_{t-1}$ is the number of wells operated by firm i at time *t-1*. $\Delta[F_{t+j} - P_t] * ShareShale$ is the percentage change in the difference between the future price at time horizon j and the spot price at time t, interacted with the share of shale wells in the total stock of wells at time $t-1$ ($\ddot{\theta}_{it-1}$).

Table A.6. Robustness: Supply elasticity - separate regressions

	(1) Pooled	(2) Conventional	(3) Shale
Δ Log Prod$_{it-1}$	-0.338***	-0.352***	-0.328***
	(-71.90)	(-75.22)	(-45.12)
Well age$_{it}$	0.001	0.000	0.001
	(1.23)	(1.23)	(0.74)
$\Delta[P_t]$	0.042	0.065*	0.011
	(1.39)	(2.38)	(0.25)
$\Delta[F_{t+1} - P_t]$	0.070	0.014	0.182**
	(1.78)	(0.43)	(2.64)
N	1202787	718142	464830
$Well FE$	Yes	Yes	Yes
$Year\,FE$	Yes	Yes	Yes
$Time\,period$	1986-2015	1986-2015	1990-2015

Note: t-statistics are in parenthesis. * Significantly different from zero at the 90% level, * 95% level, *** 99% level. The dependent variable is the monthly percentage change in the production by well i. Well age$_{it}$ is the number of months since the first production month for well i at time t. $\Delta[P_t]$ is the percentage change in the spot price of oil, while the spot future spread $\Delta[F_{t+j} - P_t]$ is the percentage change in the difference between the future price at time horizon j and the spot price at time t.

Table A.7. Robustness: Firm panel

	(Excl<20%)	(Excl shale=1)	(Shale-spec. shocks)	(Post-2008)
Dep. var: Δ Wells$_{it}$				
Δ Wells$_{it-1}$	-0.472***	-0.463***	-0.471***	-0.460***
	(-23.02)	(-21.06)	(-23.04)	(-18.71)
Stock of wells$_{t-1}$	-0.00492	-0.00319	-0.0049	-0.0481***
	(-1.53)	(-0.91)	(-1.54)	(-4.78)
$[F_{t+1} - P_t] * ShareShale$	0.175*	0.312*	0.120**	0.284**
	(2.24)	(1.99)	(2.24)	(2.09)
$[F_{t+6} - P_t] * ShareShale$	-0.340**	-0.612**	-0.23***	-1.135**
	(-2.62)	(-2.18)	(-2.71)	(-2.93)
$Time\,FE$	Yes	Yes	Yes	Yes
$Firm\,FE$	Yes	Yes	Yes	Yes
$Firms$	221	325	353	112
N	41554	53061	57187	8353
$Time\,period$	1986-2015	1986-2015	1986-2015	2009-2015

Note: t-statistics are in parenthesis. * Significantly different from zero at the 90% level, * 95% level, *** 99% level. The dependent variable is the monthly percentage change in the number of wells completed by firm i. Stock of wells$_{t-1}$ is the number of wells operated by firm i at time *t-1*. $[F_{t+j} - P_t] * ShareShale$ is the percentage change in the difference between the future price at time horizon j and the spot price at time t, interacted with the share of shale wells in the total stock of wells.

Table A.8. Hotelling: Response to net oil price changes

	(1)	(2)	(3)
Dep. var: Δ Log Prod$_{it}$			
Δ Log Prod$_{t-1}$	-0.338***	-0.338***	-0.338***
	(-71.63)	(-71.53)	(-71.54)
Well age$_{it}$	0.001	0.001	0.001
	(1.39)	(1.30)	(1.25)
$\Delta[P_t]$	0.051	0.044	0.041
	(1.69)	(1.27)	(1.05)
$\Delta[P_t] * Shale$	-0.054	-0.059	-0.056
	(-1.21)	(-1.11)	(-0.96)
$\Delta[F_{t+1} - P_t]$	0.012	0.005	0.002
	(0.35)	(0.14)	(0.06)
$\Delta[F_{t+1} - P_t] * Shale$	0.153**	0.164**	0.18***
	(3.03)	(3.20)	(3.41)
Net $\Delta[F_{t+3} - P_t]$	0.004		
	(0.01)		
Net $\Delta[F_{t+3} - P_t] * Shale$	-1.456***		
	(-3.30)		
Net $\Delta[F_{t+6} - P_t]$		0.043	
		(0.23)	
Net $\Delta[F_{t+6} - P_t] * Shale$		-0.85**	
		(-2.82)	
Net $\Delta[F_{t+12} - P_t]$			0.0321
			(0.26)
Net $\Delta[F_{t+12} - P_t] * Shale$			-0.49**
			(-2.24)
$WellFE$	Yes	Yes	Yes
$Year\,FE$	Yes	Yes	Yes
$No\,of\,Wells$	15,259	15,155	15,155
N	1,202,787	1,188,227	1,188,227
$Time\,period$	1986-2015	1986-2015	1986-2015

Note: t-statistics are in parenthesis. * Significantly different from zero at the 90% level, * 95% level, *** 99% level. The dependent variable is the monthly percentage change in the production by well i. Well age$_{it}$ is the number of months since the first production month for well i at time t. $\Delta[P_t]$ is the percentage change in the spot price of oil, while the spot future spreads $\Delta[F_{t+j} - P_t]$ is the percentage change in the difference between the future price at time horizon j and the spot price at time t.

Figure A.1. Spiky production profiles for shale wells

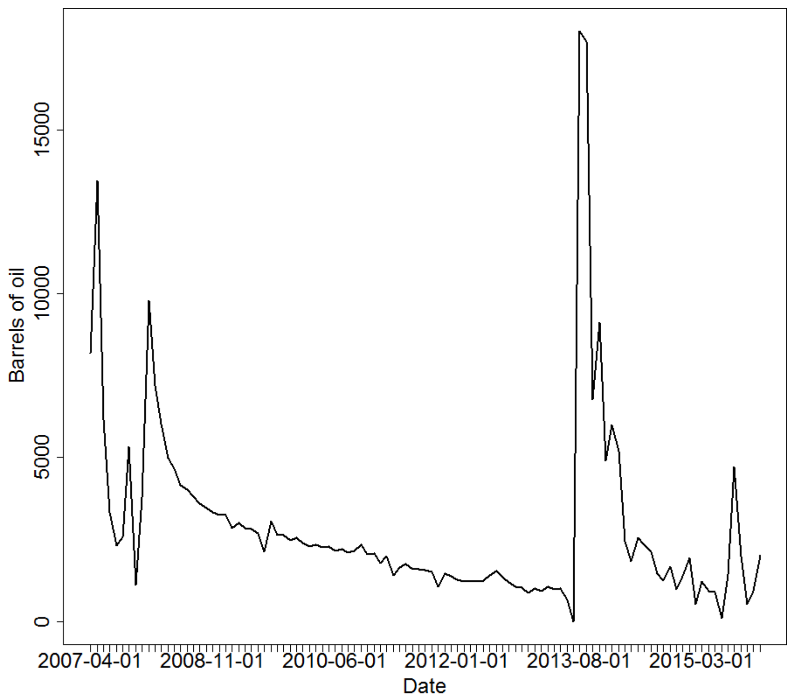

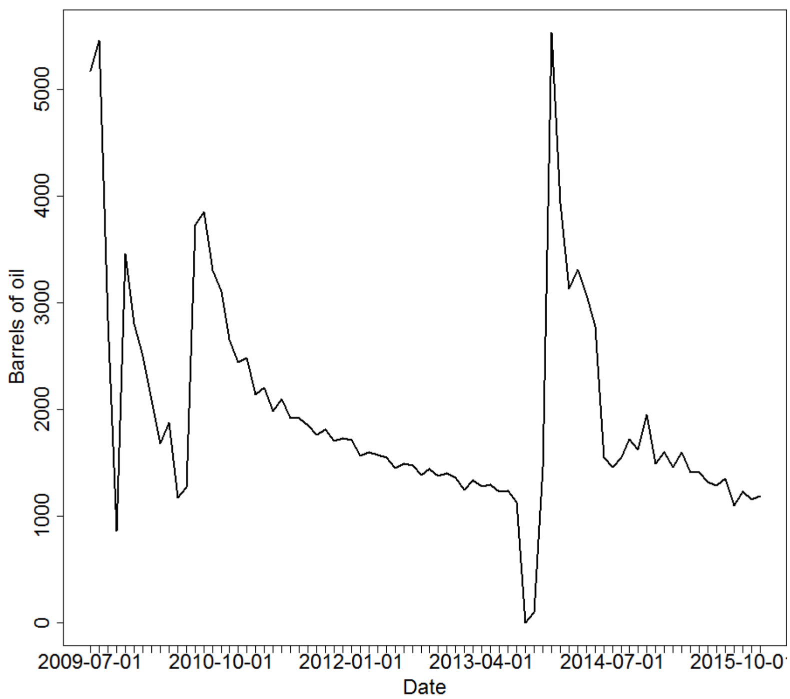

Note: The plots show production profiles for two randomly chosen shale wells in North Dakota, operated by Continental Resources inc (top graph) and Marathon oil company (bottom graph), and exhibit the spiky pattern that typically arises from repeated fracking.

Figure A.2. Share of total production and well age

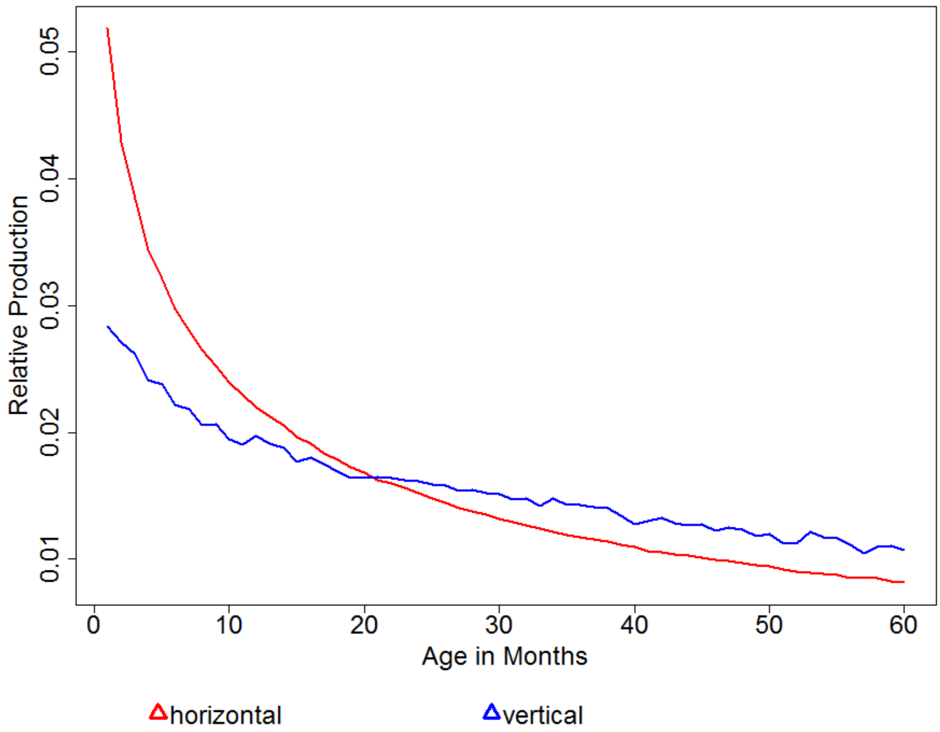

Note: The curves show the average fraction of well output occurring in each month proportional to the aggregate production over the first 60 month on the y-axis against the well lifespan measured in months on the x-axis, for the two types of well technology. The red curve shows the shale wells and the blue curve the conventional. The area under each curve is a density and sums to one.

Figure A.3. Well productivity across geographies

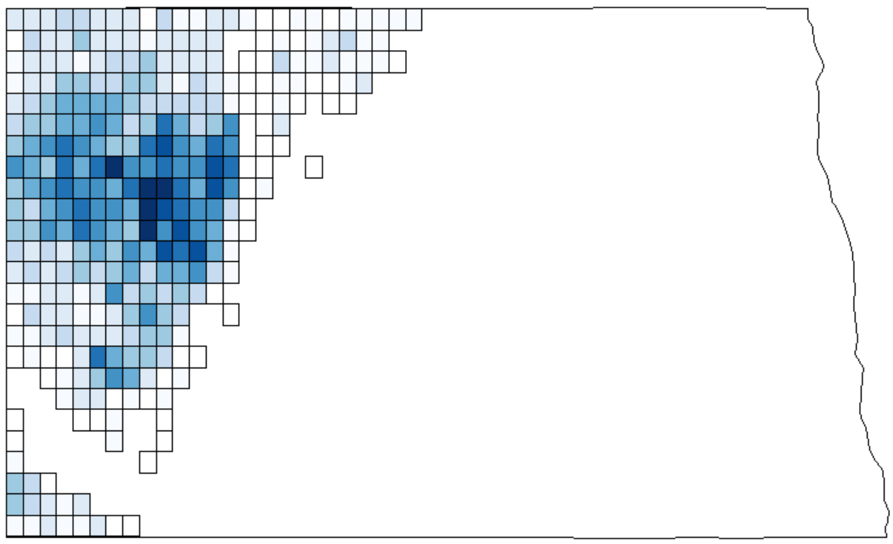

Note: The Figure shows the initial production rate averaged by each quadratic geographic unit. Darker means more productive. The geographic units are constructed using a 25 time 25 array raster covering all producing wells in North Dakota.

Figure A.4. Conventional completions and 12 month spot future spread, 1983-2015

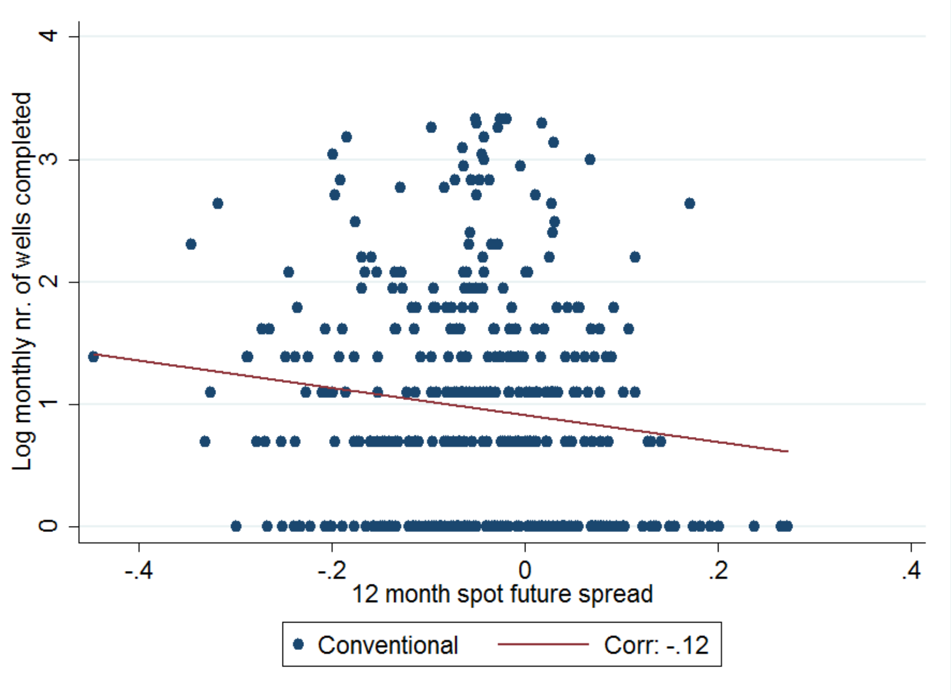

Note: The Figure shows the correlation between the expected change in the price of crude oil (spread between twelve-month future contract and the spot contract) and the number of conventional oil wells completed each month (corr.coef = -.12) from 12.1983 to 11.2015.

Figure A.5. Transition path to shale technology for 20 firms

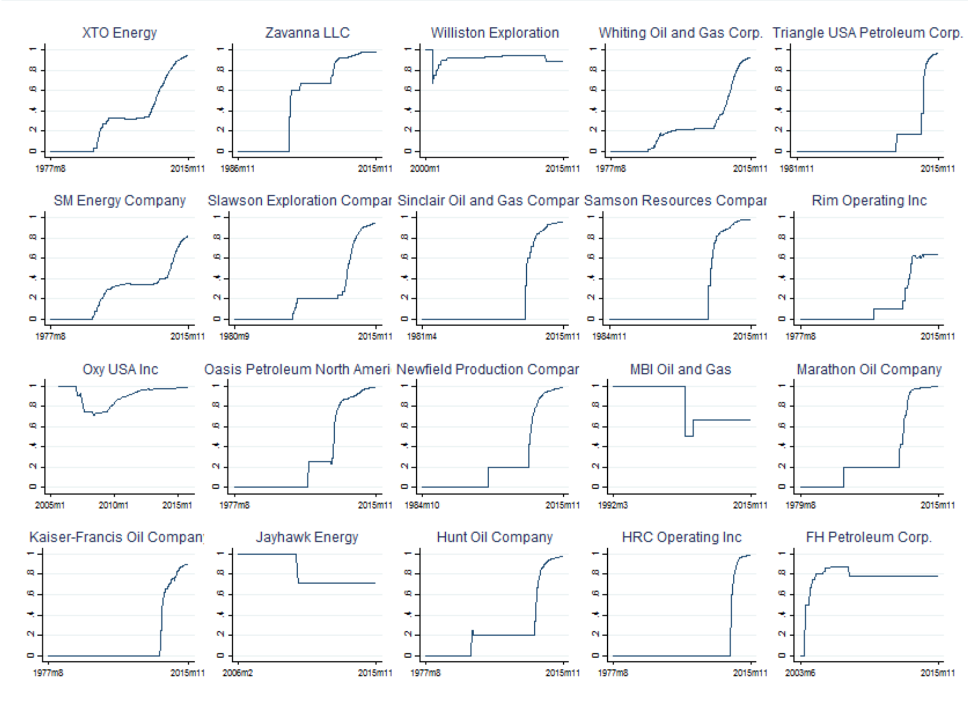

Note: The plots show the evolution of the share of shale wells in the total stock of wells, θ_i, for 20 firms for the sample time horizon.