

NBER WORKING PAPER SERIES

THE UNCONVENTIONAL OIL SUPPLY BOOM:
AGGREGATE PRICE RESPONSE FROM MICRODATA

Richard G. Newell
Brian C. Prest

Working Paper 23973
<http://www.nber.org/papers/w23973>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2017

We are grateful to Drillinginfo for data used in this study. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Richard G. Newell and Brian C. Prest. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Unconventional Oil Supply Boom: Aggregate Price Response from Microdata
Richard G. Newell and Brian C. Prest
NBER Working Paper No. 23973
October 2017
JEL No. D24,L71,Q41

ABSTRACT

We analyze the price responsiveness of onshore oil supply from conventional versus new unconventional "tight" formations in the United States. We separately analyze three key stages of oil production: drilling wells, completing wells, and production from completed wells. We find that the important margin is drilling investment. We estimate drilling responses of approximately 1.6 percent for tight oil and 1.2 percent for conventional oil per 1 percent change in oil prices. In addition, tight oil wells produce about 4.6 times more oil compared to conventional ones. Together, the long-run price responsiveness of supply is about 6 times larger for tight oil on a per well basis, and about 9 times larger when also accounting for the rise in unconventional-directed drilling. Based on our estimates derived from microdata, we conduct aggregate simulations of incremental oil supply at different time frames and price levels. The simulations show that the U.S. supply response is much larger now due to the shale revolution. Given a price rise to \$80 per barrel, U.S. oil production could rise by 0.5 million barrels per day in 6 months, 1.2 million in 1 year, 2 million in 2 years, and 3 million in 5 years. Nonetheless, it takes many months before a substantial portion of the full supply response is online, longer than the 30 to 90 days typically associated with the role of "swing producer" such as Saudi Arabia.

Richard G. Newell
Duke University
Box 90328
Durham, NC 27708
and NBER
newell@rff.org

Brian C. Prest
Duke University, Box 90328
Durham, NC 27708
brian.prest@duke.edu

Crude oil is the largest commodity market and the shale revolution has dramatically altered U.S. oil supply and markets in recent years. New technological developments such as seismic imaging, hydraulic fracturing, and horizontal drilling have unlocked vast “tight oil” reserves previously thought to be nonviable. This has led to the largest and most rapid surge in oil production in U.S. history, amounting to several million barrels per day (BPD) of additional supply in just a few years (see Kilian 2017*a* for an overview of the impact on oil and gasoline prices). This dramatic expansion in the United States’ role in oil supply suggests a larger ability of oil production to respond to price fluctuations, increasing supply responsiveness. Along with one of the most significant downturns in global oil prices and questions surrounding OPEC’s interest in acting as a market stabilizer, the tight oil boom has also raised questions about whether U.S. unconventional oil might play the role of “swing producer”. In this paper, we apply the methods from Newell, Prest and Vissing (2016) to investigate the relative price responsiveness of unconventional and conventional oil drilling in the United States to understand the degree to which these supply dynamics have fundamentally shifted.

We estimate the differences in price responsiveness for oil wells using a detailed dataset on 164,000 oil wells in the five major oil-producing states of Texas, North Dakota, California, Oklahoma, and Colorado. We estimate the price responsiveness at three key stages of production: drilling, spud-to-production time, and production from existing wells. We find that the important margin for the price response is drilling activity, estimating drilling responses of 1.6 for unconventional oil wells and 1.2 for conventional wells. This price response of oil drilling is substantially larger than the responses estimated for gas drilling in Newell, Prest and Vissing (2016) for both conventional and unconventional.

Moreover, we find that the much higher productivity of unconventional oil wells (which is about 9 times larger initially and 4.6 times larger cumulatively) magnifies this greater drilling responsiveness many times over. We conduct simulations to combine the different stages to show how the rise of unconventional drilling has affected the responsiveness of U.S. oil supply. The larger estimated drilling response combined with the larger amount of oil produced per well leads to an estimated 6-fold larger long-run price response from unconventional oil wells on a per-well basis. Further accounting for recent changes in the composition of drilling activity (unconventional versus conventional), changes in market prices, and rising per-well productivity makes this difference even larger, implying a 9-fold larger price response of U.S. supply compared to the pre-shale era.

We use our simulations to approximate an aggregate U.S. oil supply curve based on our estimates derived from microdata. We run simulations calibrated to the market situation as of 2015, estimating incremental oil production at different price levels (\$50 to \$80 per

barrel) and time frames (6 months, 1 year, 2, years, and 5 years). The simulations suggest that given a price rise to \$80 per barrel, U.S. production could rise by 0.5 million barrels per day in 6 months, 1.2 million in 1 year, 2 million in 2 years, and 3 million in 5 years. These magnitudes are significant in the context of the global market, suggesting a significantly larger role for the United States as an incremental producer. However, the time needed to drill and complete wells imply that the production response takes longer than is typically considered for a “swing producer”, which has typically been taken to mean a supplier that can increase oil production substantially (say, 1 million barrels per day) in a short period of time (within 30 to 90 days).

I Literature

This paper contributes to a growing line of research studying the shale revolution (e.g., Joskow 2013; Covert 2015; Kilian 2017*a,b*) and price formation in oil markets in general (Hamilton 2009; Kilian 2009; Anderson, Kellogg and Salant 2017; Baumeister and Kilian 2016*a,b*).

The literature on price responsiveness of oil supply often compares results for both the short-run and long-run, typically finding smaller short-run supply responses, as expected. Nevertheless, Baumeister and Hamilton (2015) find evidence for a small, positive short-run supply elasticity. Papers analyzing oil extraction elasticities include Griffin (1985); Hogan (1989); Jones (1990); Dahl and Yücel (1991); Ramcharran (2002); and Güntner (2014). Much of this work pre-dates the shale revolution, and even the more recent literature does not distinguish between unconventional and conventional supply.

A few recent studies have touched upon U.S. oil supply elasticities. Anderson, Kellogg and Salant (2017) estimate drilling elasticities in the state of Texas using time series data during the 1990-2008 period, finding an elasticity of 0.6. Smith (2017) uses sequential sampling combined with drilling cost estimates to estimate economic recoverable resources in the Bakken shale formation at different price levels. Collins and Medlock (2017) provides a qualitative overview of the ability of shale producers to scale-up activity, suggesting that high and rising productivity is a major factor. Our paper applies and builds on the methodology established in Newell, Prest and Vissing (2016), which focused on natural gas supply, to estimate the U.S. oil supply response at different stages of the production process. We use more recent data than Anderson, Kellogg and Salant (2017) (which focuses on 1990-2008) to incorporate changes due to the rise of horizontal drilling and hydraulic fracturing for oil extraction, which largely took place during the 2010-2015 period.

II Data

We use well-level data assembled by Drillinginfo, a company that provides information services on upstream oil and natural gas activity. We use Drillinginfo data¹ on oil wells in five states that collectively account for nearly 60 percent of total U.S. onshore and offshore oil production: Texas (32 percent), North Dakota (11 percent), California (7 percent), Oklahoma (4 percent), and Colorado (3 percent).² These states account for an even larger share of U.S. drilling activity at 74 percent in 2014. These states also account for virtually all tight oil plays in the United States, including the Permian, Eagle Ford, Niobrara, Bakken, and Monterey formations. After cleaning the data (discussed below), the wells in our dataset account for nearly half of all U.S. oil production and nearly two-thirds of U.S. drilling activity in 2014, the last full year of our data.³

We drop duplicated observations and observations with missing or invalid dates.⁴ We focus on oil wells and do not include gas wells in this analysis.⁵ While some gas wells do co-produce oil, this share of oil supply is small. We compute the length of horizontal well “laterals” using the geodesic distance between the well’s surface hole and bottom hole. We consider a well to be unconventional if it was drilled horizontally; otherwise, we consider the well to be conventional.⁶

Unless otherwise noted, we use the average of the next 12 months of futures prices for oil (WTI) and gas (Henry Hub) from Bloomberg. Each price is the average of daily prices and adjusted to 2014 dollars with the CPI All Urban Consumer (All Items) index.⁷

¹We downloaded the Drillinginfo dataset on August 23, 2016.

²Percentages represent averages during the shale oil era of 2010 to 2016 based on https://www.eia.gov/dnav/pet/pet_crd_crpdn_adc_mbbldpd_m.htm and https://www.eia.gov/energy_in_brief/article/shale_in_the_united_states.cfm.

³In particular, the wells produced on average 4 million barrels per day in 2014, the last full year of our data. According to EIA, total U.S. oil production averaged 8.8 million barrels per day that year (https://www.eia.gov/dnav/pet/pet_crd_crpdn_adc_mbbldpd_a.htm).

⁴As described in Newell, Prest and Vissing (2016), such inconsistencies are due to the way well data is stored and updated.

⁵The definition of a “gas” well versus an “oil” well is determined by regulatory standards that are designed to ensure adequate well spacing.

⁶This is a slight departure from the method in Newell, Prest and Vissing (2016), where we also used reservoir information to help classify gas wells and also treated certain directionally-drilled gas wells in Texas as unconventional. We did this because some directionally-drilled oil wells outside of Texas are clearly conventional based on when they were drilled, their reservoirs, their productivities, and their type curves. In particular, in California and Colorado the type curves of directionally-drilled oil wells more closely resemble type curves of vertical oil wells, compared to horizontal ones. For this reason, we classify directionally-drilled wells as conventional. In the other three states, there are relatively few directionally-drilled wells, so this small change of classification method does not matter for those wells.

⁷Using futures prices as a measure of price expectations is a shortcut to obtain price expectations. This is based on conversations with industry operators regarding how they generate their price expectations. Baumeister and Kilian (2016b) show that futures prices and expected spot prices can diverge in the presence of a

The dataset describes both static and dynamic characteristics of each well. The static characteristics include each well’s important dates (spud and first production dates), location, drilling direction, and reservoir, among other features. The dynamic characteristics include a monthly time series of each well’s oil and gas production. In all states except Texas, oil production is measured at the well level. In Texas, oil production is measured at the lease level and Drillinginfo allocates oil production to individual wells using well test data.

The cleaned dataset includes approximately 164,000 onshore oil wells drilled between 2000 and 2015. As in Newell, Prest and Vissing (2016), we only include unconventional wells drilled during the shale revolution period of 2005-2015. The map in Figure 1 illustrates the location of the wells in our data along with selected shale plays.

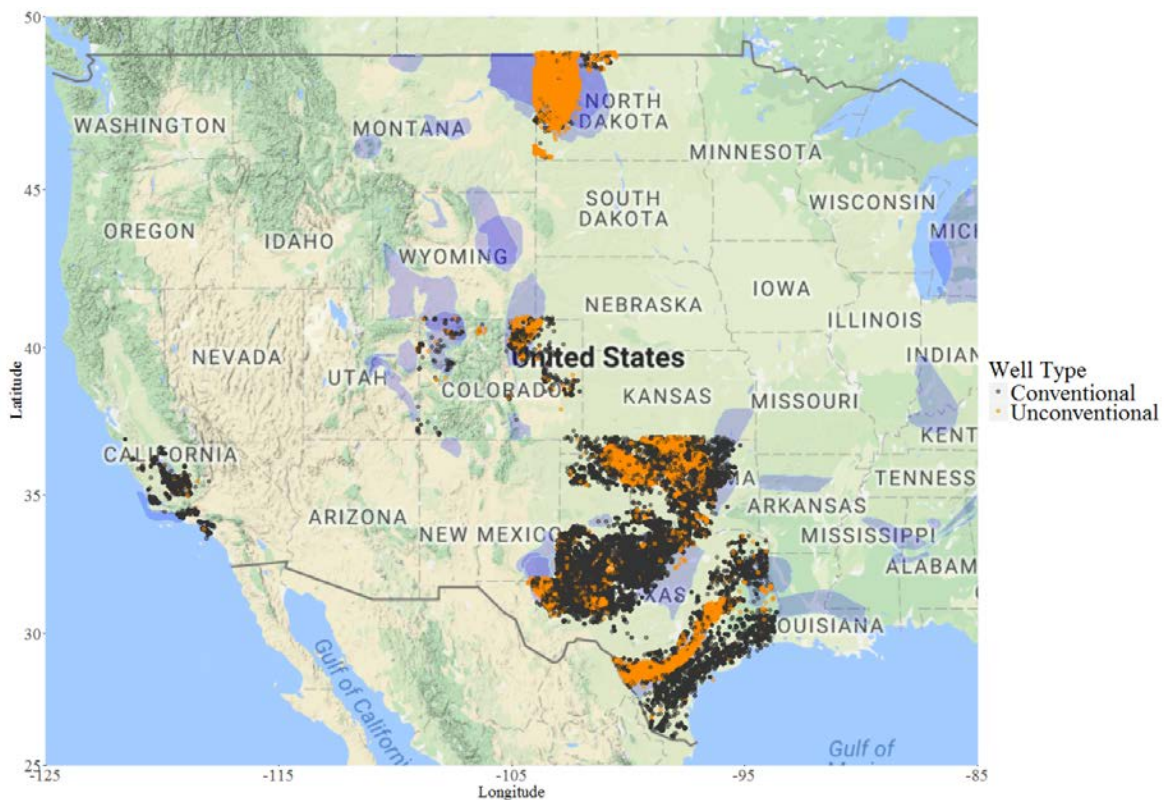


Figure 1: Location of Oil Wells in Data by Well Type and Selected Shale Plays

Sources: Well locations are from Drillinginfo. Map is from Google via the ggmap package for R developed by Kahle and Wickham (2013). The indication of shale formations is based on EIA’s shapefile for low permeability oil and gas play boundaries in the Lower 48 States, available at <https://www.eia.gov/maps/maps.htm#geodata>.

time-varying risk premium, particularly for long-dated oil futures contracts. The combination of discounting and approximately exponential decline curves somewhat mutes the impact of this distinction on the expected net present value of revenues.

Figure 2 shows the counts of wells spudded from 2000.Q1 to 2015.Q3 by well type, along with oil and natural gas prices in barrels of oil equivalent (boe) units. One can see a clear relationship between drilling and prices for both well types, particularly during 2007-2009 and 2014-2015. We analyze this relationship further in section III.A.

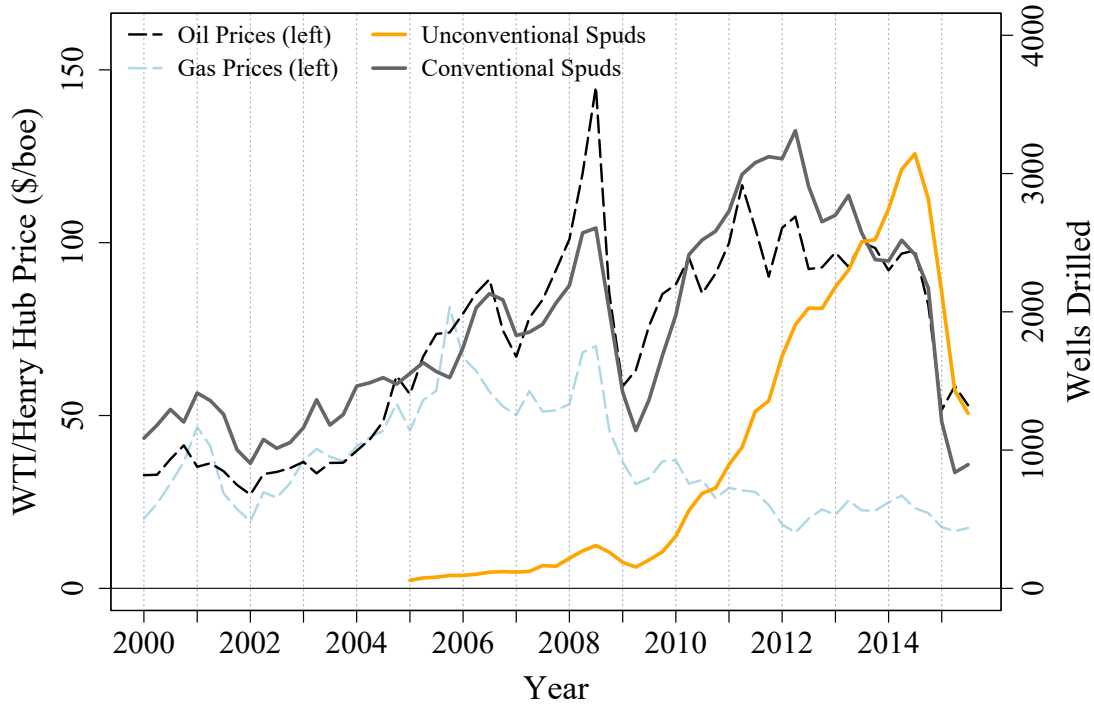


Figure 2: Number of Spuds by Well Type (right axis) and Oil & Gas Prices (left axis), 2000-2015, Quarterly

Sources: Authors' calculations based on data from Drillinginfo and Bloomberg.

Table 1 reports the summary statistics for the wells and price data. Unconventional oil wells produce much more oil per well than conventional wells do, particularly during their first month. Average initial productivity (IP) is 348 barrels per day (BPD), approximately 9 times the average for conventional wells (40 BPD) in this period.⁸ In addition, the productivity of unconventional wells has been expanding steadily over time, with more recent wells having average productivity closer to 400 BPD. By contrast, average productivity of conventional wells has remained flat for the past two decades.⁹

⁸The t-statistic for this difference is -184.43. Initial production is measured as production during its first full month of production, meaning the second calendar month during which production is reported. It is standard to focus on the second month because a well is typically only producing for a fraction of its first calendar month.

⁹See appendix for graphs showing these trends.

While average initial production is approximately 9 times larger for unconventional wells, they have steeper decline rates, implying that the productivity advantage shrinks over time. As a result, they produce only about 6.5 times much on average over the first 12 months (63,253 barrels versus 9,689 barrels).¹⁰ Over the longer-run, the cumulative unconventional production advantage is about 4.6.

Unconventional wells are also less variable, as evidenced by the relative similarity of their mean and median productivity, as well as its lower coefficient of variation (1.0 versus 1.8).

Table 1: Summary Statistics

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
	Conventional			Unconventional		
	Mean	Median	Std. Dev.	Mean	Median	Std. Dev.
Well Data						
Initial Oil Production (first full month, barrels per day)	40	22	71	348	270	351
First 12 Months' Total Oil Production (barrels)	9,689	5,625	16,123	63,253	51,596	55,106
Oil 3-Month Decline Rate (%)	49.1	46.4	27.6	53.5	52.1	23.2
Oil 12-Month Decline Rate (%)	72.6	76.2	22.9	78.6	80.3	16.4
Oil 24-Month Decline Rate (%)	81.8	86.8	19.1	86.6	88.4	12.3
<hr/>						
Initial Gas Production (first full month, mcf per day)	73	10	419	525	309	898
First 12 Months' Total Gas Production (mcf)	19,666	3,966	81,735	117,319	78,118	161,268
Gas 3-Month Decline Rate (%)	51.7	48.0	29.9	52.1	49.0	27.2
Gas 12-Month Decline Rate (%)	74.1	77.6	23.9	75.2	77.4	21.1
Gas 24-Month Decline Rate (%)	82.8	88.8	19.7	83.7	86.8	16.4
<hr/>						
Horizontal Well Length (ft)				5,532	5,346	2,879
Total Vertical Depth (ft)	5,221	4,958	3,953	15,079	15,303	4,265
Months Between Spud Date and First Production	2.67	2.00	3.06	4.10	3.00	2.87
Number of Wells		118,774			44,918	
<hr/>						
Price Data (Monthly, 2000-2015)						
	Mean	Median	Std. Dev.			
WTI Oil Price - Prompt Month Future (\$/barrel)	\$70.90	\$73.12	\$27.17			
WTI Oil Price - 12-Month Future (\$/barrel)	\$71.15	\$76.13	\$27.78			
Henry Hub Natural Gas Price - Prompt Month Future (\$/MMBTU)	\$5.95	\$5.02	\$2.71			
Henry Hub Natural Gas Price - 12-Month Future (\$/MMBTU)	\$6.31	\$5.63	\$2.64			

Sources: Authors' calculations based on data from Drillinginfo and Bloomberg.

However, unconventional wells take somewhat longer to begin production (4.1 months versus 2.7 on average) after they have been “spudded” (that is, after drilling has begun). We attribute this to the time needed to drill the longer wellbores and hydraulically fracture them. We explore this spud-to-production time profile in more detail in section III.B.

III Models and Results

We divide our analysis into three stages, as in Newell, Prest and Vissing (2016): drilling (or “spudding”) activity (section III.A), spud-to-production time (section III.B), and production from existing wells (section III.C). Then in section III.D we conduct simulations that explicitly integrate the modeling of these three stages.

¹⁰The t-statistic for this difference is -202.75.

III.A Stage 1: Commence Drilling (Spud) a Well

III.A.1 Drilling Estimation Method

To estimate the relationship between drilling activity and prices, we estimate the following specification of drilling activity as a function of estimated oil and gas revenues:¹¹

$$\Delta \ln(w_{j,t}) = \beta_0 + \sum_{l=0}^L [\beta_{1,j,l} \Delta \ln(\tilde{p}_{oil,t-l} \tilde{q}_{oil,j,t-l}) + \beta_{2,j,l} \Delta \ln(\tilde{p}_{gas,t-l} \tilde{q}_{gas,j,t-l})] + \varepsilon_t \quad (1)$$

where $w_{j,t}$ is the count of oil spuds of well type j (conventional or unconventional) in quarter t and Δ indicates estimation in first differences. Estimated oil and gas revenues are a well's productivity (denoted $\tilde{q}_{oil,t}$ and $\tilde{q}_{gas,t}$ respectively) multiplied by the price ($\tilde{p}_{oil,t}$ and $\tilde{p}_{gas,t}$) of the well's output. As described in Newell, Prest and Vissing (2016), we use revenues as the explanatory variable, rather than prices, to reflect the fact that the returns to drilling a well are given by the total value of its production, rather than the price of a single unit of its output (however we also estimate using simply prices as a sensitivity). The parameters of interest are $\beta_{1,j,l}$, representing the l -lagged drilling response. We use $L = 3$ quarterly lags of the revenue variables to account for the fact that drilling decisions are often made months before the drilling rig arrives on site. The cumulative drilling response with respect to oil prices¹² is given by $\sum_{l=0}^L \beta_{1,j,l}$.

For the price variables, we use the simple average of the next 12 months of WTI and Henry Hub futures prices adjusted for inflation to 2014 dollars. For the productivity variables ($\tilde{q}_{oil,j,t}$ and $\tilde{q}_{gas,j,t}$), we use a two-quarter moving average of initial production, by well type.¹³

We consider the potential endogeneity of oil prices, owing to unobserved cost shocks that could both increase drilling activity and reduce oil prices, creating a negative covariance term that could bias estimates of the price response toward zero. The need to instrument for price endogeneity is widely recognized in the literature estimating demand elasticities for fossil fuels (Davis and Kilian 2011 and Coglianese et al. 2017), suggesting that similar issues may also arise when estimating supply elasticities.¹⁴ In particular, the shale revolution has arguably had some effect on oil prices in recent years, raising concerns about the endogeneity of oil prices. Indeed, Kilian (2016), Kilian (2017a), and Kilian (2017b) find that the shale

¹¹We estimate the relationship in first differences to make the revenue and drilling series stationary. For an explanation of this specification, see Newell, Prest and Vissing (2016).

¹²We do not distinguish between "revenue" versus "price" responses because they are equivalent, holding productivity constant, as in the case of an exogenous price shock. This owes to the following equality: $\beta_{1,j,l} \ln(\tilde{p}_{oil,t-l} \tilde{q}_{oil,j,t-l}) = \beta_{1,j,l} \ln(\tilde{p}_{oil,t-l}) + \beta_{1,j,l} \ln(\tilde{q}_{oil,j,t-l})$.

¹³These data are presented in the appendix.

¹⁴Studies of U.S. oil supply elasticities often do not instrument for oil prices based on the historically plausible argument that incremental production from the United States is small relative to the global oil market. This argument requires that drilling cost shocks affecting oil prices are from foreign (not U.S.) supply. The rise of the shale revolution makes this argument less tenable, however.

revolution had some effect on WTI prices during 2011 through mid-2014, although the effect is negligible thereafter. On the other hand, Baumeister and Kilian (2016a) demonstrate that fluctuations in oil prices are mostly driven by demand factors, suggesting a lesser need for addressing price endogeneity in this setting.

We instrument for the potential endogeneity of oil prices using the Commodity Research Bureau (CRB) Raw Industrial Commodity Index. The CRB Index was collected from Bloomberg L.P.¹⁵ Baumeister and Kilian (2012) show that this index is a good predictor of oil prices. Barsky and Kilian (2001) noted that oil prices move together with the prices of other industrial commodities, suggesting that these co-movements are primarily driven by demand factors.

The sample period spans 2000.Q1 to 2015.Q3 for conventional wells and 2005.Q1 to 2015.Q3 for unconventional wells;¹⁶ the unit of observation is one quarter. All standard errors are Newey-West.

III.A.2 Drilling Estimation Results

Table 2 presents our results from estimation of equation (1). The top four rows show the response of drilling with respect to changes in oil prices/revenues at $l = 0, 1, 2, 3$ lags, and the cumulative price response is shown in the middle of the table along with its standard error. For each type of well (conventional and unconventional), the cumulative response is generally consistent across specifications, with an estimate of about 0.9-1.3 for conventional wells and an estimate of about 1.2-1.9 for unconventional wells.

These are large responses, consistent with the strong adjustments exhibited in Figure 2. For example, consider the steep price decline in from 2014.Q3 to 2015.Q3, the last 12 months of our sample. In that time, the WTI price fell by 46 percent (from \$98 to \$53). In that same time, U.S. drilling activity fell by approximately 60 percent. A back-of-the-envelope calculation from that experience alone would suggest an elasticity of about 1.3 ($\approx \frac{-0.60}{-0.46}$), in the middle of our range.

For each type of well, three specifications are shown. The first specification is our baseline specification, which uses oil and gas revenues as the explanatory variables and instruments for the endogeneity of oil prices.¹⁷ The second specification is a sensitivity in which

¹⁵The CRB index replaces the instrument of copper prices used in a previous draft because the CRB index proved to be a stronger instrument.

¹⁶Expanding the time period for unconventional wells back to 2000 would be inappropriate because unconventional drilling was virtually non-existent in this period.

¹⁷We only instrument for contemporaneous oil prices, treating lagged prices as exogenous. As a sensitivity, we also estimated our specification treating lagged oil prices as endogenous by conducting a single first stage regression and taking its lags. The estimated responses were largely unchanged, at 1.1 and 1.5 for conventional and unconventional wells respectively, compared to our benchmark estimates of 1.2 and 1.6.

we use oil and gas prices, instead of revenues.¹⁸ The first stages for these regressions can be found in the appendix. Finally, we show the results from an OLS regression, where all regressors are treated as exogenous.

The first three columns show the results of the estimation for conventional wells. Estimation of our base specification shown in column (1) finds a cumulative drilling response of 1.2. The corresponding value for unconventional wells is 1.6, shown in column (4). While these estimates are each statistically different from zero, their confidence intervals overlap.

Columns (2) and (5) show the effect of using oil and gas prices, rather than revenues, which increases the cumulative price responses to 1.3 and 1.9 for conventional and unconventional wells, respectively. Nonetheless, we prefer our base specification using revenues, which more accurately represent drilling incentives.

Columns (3) and (6) show the results of re-estimating equation (1) by ordinary least squares, without the instruments (that is, assuming oil prices are exogenous to U.S. drilling activity). The price response estimates of 0.9 and 1.2 are somewhat smaller than the instrumental variables (IV) estimates of 1.2 and 1.6 from columns (1) and (4). However, the Wu-Hausman test fails to reject the null of “no price endogeneity” for the unconventional equation. This is consistent with Baumeister and Kilian (2016a) and Kilian (2016), who argue that the shale revolution has played a relatively smaller role in price formation compared to demand factors, which are the primary determinants of oil price fluctuations. On the other hand, the Wu-Hausman test does reject the null for the conventional equations.

While oil drilling appears to respond strongly to oil price movements, there is a weaker response to natural gas prices. For conventional drilling, the estimated response to gas prices is negligible. For unconventional drilling, there is a moderate negative relationship, consistent with a substitution effect. Indeed, industry experts have noted the shift of rigs from oil-directed drilling towards gas-directed drilling in recent years, owing to falling natural gas prices.

In summary, both conventional and unconventional drilling appears to have a response somewhat larger than 1. This estimate is larger than the estimates from the recent Anderson, Kellogg and Salant (2017) study, which estimates a response of approximately 0.6 for drilling in Texas during 1990-2008. The specification in that study only allows drilling to respond the most recent three months of price changes (current-month changes plus two monthly lags). The analogous figure in our estimation would be the response to shocks during the contemporaneous quarter (also a three-month period) for conventional wells (since their 1990-2008 sample predates the boom in tight oil). We estimate this to be a comparable

¹⁸As a sensitivity, we also estimated a specification where prices and productivity are allowed to enter separately. The estimated price responses in that specification were virtually identical to those with just prices alone.

0.67 (seen in the first row, first column of Table 2). The rest of our estimated response owes primarily to the additional drilling response to price changes during the first quarterly lag.

In the next section, we consider the responsiveness of the time from spud to first production from wells that are drilled.

Table 2: Drilling Estimation Results

	Conventional Wells			Unconventional Wells		
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta\text{Log}(\text{Oil Revenues or Prices})$	0.67 (0.12)	0.75 (0.15)	0.45 (0.07)	1.07 (0.26)	0.81 (0.17)	0.59 (0.10)
$\Delta\text{Log}(\text{Oil Revenues or Prices}), 1 \text{ Lag}$	0.46 (0.10)	0.52 (0.12)	0.43 (0.12)	0.37 (0.17)	0.56 (0.16)	0.36 (0.18)
$\Delta\text{Log}(\text{Oil Revenues or Prices}), 2 \text{ Lags}$	0.08 (0.09)	0.17 (0.07)	0.04 (0.07)	0.18 (0.07)	0.47 (0.12)	0.18 (0.10)
$\Delta\text{Log}(\text{Oil Revenues or Prices}), 3 \text{ Lags}$	-0.02 (0.07)	-0.12 (0.10)	0.01 (0.08)	0.01 (0.06)	0.11 (0.12)	0.02 (0.09)
$\Delta\text{Log}(\text{Gas Revenues or Prices})$	0.09 (0.09)	-0.03 (0.08)	0.14 (0.09)	0.01 (0.08)	-0.10 (0.12)	0.12 (0.07)
$\Delta\text{Log}(\text{Gas Revenues or Prices}), 1 \text{ Lag}$	-0.07 (0.05)	-0.10 (0.08)	-0.05 (0.05)	-0.34 (0.11)	-0.32 (0.14)	-0.20 (0.05)
$\Delta\text{Log}(\text{Gas Revenues or Prices}), 2 \text{ Lags}$	0.01 (0.06)	0.05 (0.04)	0.05 (0.05)	-0.04 (0.13)	-0.11 (0.10)	0.04 (0.08)
$\Delta\text{Log}(\text{Gas Revenues or Prices}), 3 \text{ Lags}$	-0.01 (0.06)	0.02 (0.06)	0.01 (0.06)	-0.07 (0.10)	-0.23 (0.10)	-0.10 (0.07)
2nd Quarter Indicator	-0.05 (0.03)	-0.03 (0.02)	-0.03 (0.02)	-0.15 (0.10)	-0.06 (0.04)	-0.05 (0.05)
3rd Quarter Indicator	-0.08 (0.02)	-0.10 (0.03)	-0.08 (0.02)	-0.08 (0.10)	0.01 (0.04)	-0.02 (0.06)
4th Quarter Indicator	-0.04 (0.04)	-0.09 (0.04)	-0.05 (0.04)	-0.07 (0.04)	-0.07 (0.03)	-0.08 (0.03)
Constant	0.03 (0.02)	0.04 (0.02)	-0.01 (0.01)	0.11 (0.04)	0.08 (0.02)	0.04 (0.02)
Cumulative Oil Price Response	1.19 (0.22)	1.33 (0.31)	0.93 (0.17)	1.63 (0.43)	1.95 (0.48)	1.18 (0.51)
Cumulative Gas Price Response	0.02 (0.23)	-0.06 (0.16)	0.12 (0.19)	-0.45 (0.24)	-0.77 (0.32)	-0.18 (0.21)
Observations	63	63	63	43	43	43
R ²	0.59	0.58	0.61	0.53	0.63	0.61
Adjusted R ²	0.51	0.49	0.55	0.36	0.49	0.52
Revenues or Prices?	Revenues	Prices	Revenues	Revenues	Prices	Revenues
Estimation Method	2SLS	2SLS	OLS	2SLS	2SLS	OLS
First-stage F statistic	54.7	58.1	na	19.1	63.1	na
Wu-Hausman p-value	0.002	0.001	na	0.119	0.299	na

Sources: Authors' calculations based on data from Drillinginfo, EIA, and Bloomberg.

III.B Stage 2: Spud-to-Production Time

III.B.1 Duration Model

As shown in Table 1, once a well is spudded, it takes a few months of drilling and completion work before it begins producing oil. In this section, we model the time from spud to initial production to estimate whether operators adjust completion speed in response to price changes. We estimate the spud-to-production time distribution using survival time models with time-varying coefficients, as in Newell, Prest and Vissing (2016). We assume that this distribution follows a gamma distribution whose mean can be shifted by changes in oil and gas prices, leading to a standard maximum likelihood estimation approach.¹⁹

A survival time model parameterized with a gamma distribution is an “accelerated failure time” (AFT) model. This means that one can interpret the explanatory variables as additively shifting an observation’s expected log “failure” time, which in our setup indicates the well has begun production. Therefore, one can interpret the coefficients on log oil and gas revenues as spud-to-production time elasticities. Economic theory would suggest that these elasticities should be negative, as price increases should lead to reductions in spud-to-production time.

We consider a well to be “at risk” for 24 months after it has been spudded, standard errors are clustered at the well level, and we use the same method of computing revenues as described in section III.A.

III.B.2 Spud-to-Production Duration Estimation Results

The estimated parameters are shown in Table 3. Columns (1) and (3) show the preferred specifications, which use only oil and gas revenues as explanatory variables.²⁰ These show small negative spud-to-production elasticities for both conventional wells (-0.02, not significant) and unconventional wells (-0.14, significant).

¹⁹The gamma distribution has two ancillary parameters that make it very flexible, allowing for many non-monotonically-shaped distributions. We also tested alternative distributions, including Weibull, exponential, Gompertz, Log-normal, and Log-logistic finding that the gamma better fit the observed distribution. The density of the gamma distribution is given by,

$$f(t) = \begin{cases} \frac{\gamma^\gamma}{\sigma t \sqrt{\gamma} \Gamma(\gamma)} \exp(z\sqrt{\gamma} - u) & \text{if } \kappa \neq 0 \\ \frac{1}{\sigma t \sqrt{2\pi}} \exp(-z^2/2) & \text{if } \kappa = 0, \end{cases}$$

where $\gamma = |\kappa|^{-2}$, $z = \text{sign}(\kappa)(\ln(t) - \mu)/\sigma$, $u = \gamma \exp(|\kappa|z)$, $\Gamma(\cdot)$ is the gamma function, and we parameterize $\mu = \mathbf{X}'_{i,j,t} \boldsymbol{\theta}_j$, where $\mathbf{X}_{i,j,t}$ contains logged oil revenues and gas revenues, plus controls in the sensitivity analyses. The ancillary parameters (σ and κ) are estimated from the data, by well type.

²⁰All specifications also have spud year fixed effects to control for secular trends in drilling and completion times.

In the other columns, we further include well depth and (for unconventional wells) lateral length, both of which are directly related to the amount of time and effort involved in drilling, fracturing, and completing the wells. Adding these controls does not strongly change the estimated elasticities, which are now -0.03 (now significant) for conventional and -0.09 to -0.14 for unconventional wells. These elasticities indicate that higher prices lead to faster times from spud to production, although they are fairly small in magnitude, suggesting a minimal price response along this margin.

Table 3: Spud-to-Production Duration Model Results

	(1)	(2)	(3)	(4)	(5)
Spud-to-Production Survival Time	Conventional Oil Wells		Unconventional Oil Wells		
Log(Oil Revenues)	-0.0187 (0.0139)	-0.0319 (0.0136)	-0.136 (0.0188)	-0.135 (0.0184)	-0.0891 (0.0193)
Log(Gas Revenues)	-0.0706 (0.0126)	-0.0634 (0.0123)	0.08 (0.0178)	0.0615 (0.0174)	0.0336 (0.0180)
Log(Vertical Depth)		0.102 (0.00258)		0.238 (0.0146)	0.454 (0.00959)
Log(Lateral Length)					0.0101 (0.00194)
Constant	1.375 (0.1008)	0.5105 (0.1013)	2.067 (0.1711)	-0.1186 (0.2170)	-2.545 (0.2021)
Gamma Density Function Parameters					
σ	0.5462 (0.0014)	0.5267 (0.0014)	0.4930 (0.0020)	0.4749 (0.0021)	0.4648 (0.0021)
κ	-0.507 (0.0058)	-0.6101 (0.0067)	0.0232 (0.0124)	-0.1574 (0.0197)	-0.1048 (0.0148)
Well-Months (N*T)	354,068	354,068	204,879	204,879	179,657
Wells (N)	96,819	96,819	41,262	41,262	35,466
Log-Likelihood	-83,021.96	-81,313.77	-29,473.40	-28,087.16	-23,272.53

Notes: Clustered standard errors in parentheses. Standard errors are clustered at the well level. Coefficients can be interpreted as elasticities of expected spud-to-completion time. All specifications have fixed effects for year of spud.

Sources: Authors' calculations based on data from Drillinginfo and Bloomberg.

As just noted, the economic magnitudes of the oil elasticities estimated here are small. To further illustrate this, Figure 4 shows the estimated distributions (using specifications (1) and (3)) under two very different price levels: \$50 per barrel and \$100 per barrel. At the much higher price of \$100, the unconventional distribution shifts only slightly to the left, and the conventional distribution barely changes at all. In short, once wells have been drilled, there appears to be little response of completion speed to price levels.

The coefficients on well depth and length are strongly positive, suggesting it takes significantly longer to complete deeper and longer wellbores, which is intuitive. We find small coefficients of mixed-signs on the price of the well's co-produced by-product (here, the price of natural gas), consistent with Newell, Prest and Vissing (2016). The gas price elasticities

range from -0.07 to +0.08. Consistent with the drilling results in the previous section, this suggests that oil completion activity does not respond strongly to natural gas prices.

The ancillary parameters of the gamma distribution are also estimated, both for unconventional and conventional wells. The fitted distributions, plotted at covariate means, are shown in Figure 3, along with kernel density estimates of the raw distribution. This shows that the estimated gamma distributions fit the true distributions very well, validating the gamma distributional assumption.




Figure 3: Estimated Spud-to-Production Time Distribution, by Well Type

Sources: Authors' calculations based on data from Drillinginfo.




Figure 4: Illustration of Oil Price Effect on Spud-to-Production Time, by Well Type

Sources: Authors' calculations based on data from Drillinginfo and Bloomberg.

III.C Stage 3: Production Profile over Time

III.C.1 Production Profile Estimation Method

Once wells are online, oil and gas flow to the surface for many years. A well's flow rate depends on subsurface pressure, meaning that wells tend to produce at their highest rates immediately, followed by a quick decline. The average production profiles for unconventional and conventional wells in our data are shown in Figure 5.²¹

For the most part, once a well is online its flow rate is determined by geology and is therefore out of the operator's control. In some circumstances, however, operators may choose to choke production or artificially stimulate it, but doing so can risk reducing how much of the underlying resource can be ultimately recovered. Altogether, economic theory combined with physical constraints suggest that the price responsiveness of production from existing wells should be small or even zero (see Anderson, Kellogg and Salant 2017).

²¹These curves were computed as described in Newell, Prest and Vissing (2016). In particular, they represent simple averages (or medians) of oil production by age of well.




Figure 5: Mean and Median Profile of Monthly Oil Production from Oil Wells

Sources: Authors’ calculations based on data from Drillinginfo.

In this section, we empirically estimate the time profile of oil production, i.e., the “decline curve” for oil wells, including the degree to which the oil production profile of existing wells is responsive to oil prices. Moreover, even if this profile is not price-responsive, the shape of the decline curve is important to understanding the time profile of the supply response. In particular, to the extent that decline curves are steep, meaning wells produce most of their output shortly after coming online, there is a tighter temporal relationship between drilling activity and realized oil production

We estimate the price responsiveness of production from existing wells by estimating the following linear fixed effects specification:

$$\ln(q_{i,oil,j,t}) = \chi_i + \eta_{oil,j} \ln(p_{oil,t}) + \eta_{gas,j} \ln(p_{gas,t}) + g_j(Age_{i,t}) + \varepsilon_{i,j,t}, \quad (2)$$

where $q_{i,oil,j,t}$ is oil production from well i of type j (unconventional or conventional) in month t . $p_{oil,t}$ and $p_{gas,t}$ are the spot oil price (WTI) and natural gas price (Henry Hub) at time t . χ_i is a well-level fixed effect to capture the overall productivity of each well. $g_j(Age_{i,t})$ is

a flexible function of well age (using polynomials or cubic splines²²) to capture the natural decline in oil production over time from wells of type j . The key parameter of interest is $\eta_{oil,j}$, which is the price elasticity of oil supply *from existing wells* of type j .

III.C.2 Production Profile Estimation Results

The results from estimating equation (2) for each well type (unconventional and conventional) are shown in Table 4. Reading the columns from left to right, the table shows the results using increasingly flexible time trends for the decline curves. Across these specifications, we generally find small estimated price elasticities for both unconventional and conventional wells. We prefer the most flexible functional form, the cubic spline, shown in column (5). The estimated elasticities under this specification are 0.12 for unconventional wells and -0.02 for conventional ones.²³ These estimates are small but nonetheless statistically significant owing to the large sample size of 9.5 million well-month observations.²⁴ As a sensitivity, we also estimated these equations in first differences, again finding negligible price responses. We present the results in levels to ease the reader's use of the coefficients for modeling decline curves.

III.D Simulation of Oil Production Response

We conduct several simulations that link together the three individual models described above: drilling activity, time from spud to production, and the time profile of realized production. These simulations demonstrate how a permanent price shock affects drilling activity and production over time according to the three models.

The first simulation is roughly calibrated to 2015 values of baseline drilling activity, but normalized for each well type to present an apples-to-apples comparison of the relative price responsiveness of the two well types. The second set of simulations considers the total price responsiveness of oil supply from conventional and unconventional sources combined, and assesses how this responsiveness has evolved over time with the changing nature of U.S. oil supply. A third set of simulations is calibrated to the 2015 market environment, but analyzes

²²The splines use knots at every 12-month interval following initial production.

²³The one specification with a substantial elasticity is for unconventional wells in column (1), with a point estimate of 0.42. However, that specification uses an inflexible linear time trend for log oil production, effectively assuming that production decline exponentially in levels. This is likely to be insufficiently flexible for shale wells, which typically decline faster than exponentially. Indeed, this is confirmed by the estimates in columns (2) through (4) that show substantial and significant positive coefficients on the quadratic age terms.

²⁴Standard errors are clustered at the well level. We drop the first month of production, which are typically partial months, as seen in the average curves depicted in Figure 5.

Table 4: Well Production Profile Fixed Effects Regressions

Dep. Var.: Log(Oil Production)	(1)	(2)	(3)	(4)	(5)
Unconventional Wells					
Log(Oil Price)	0.42 (0.013)	0.15 (0.012)	0.09 (0.012)	0.11 (0.012)	0.12 (0.012)
Log(Gas Price)	-0.201 (0.02)	-0.04 (0.019)	-0.03 (0.019)	-0.03 (0.019)	-0.04 (0.019)
Well Age (months)	-0.027 (0.000312)	-0.06 (0.000546)	-0.093 (0.000894)	-0.122 (0.001265)	
Well Age ² (months)		0.00039 (0.000006)	0.00131 (0.000023)	0.00269 (0.000052)	
Well Age ³ (months)			-0.000006 (0.0000002)	-0.000028 (0.0000008)	
Well Age ⁴ (months)				0.0000001 (0.000000004)	
Conventional Wells					
Log(Oil Price)	0.004 (0.004)	0.01 (0.004)	-0.03 (0.004)	-0.03 (0.004)	-0.02 (0.004)
Log(Gas Price)	-0.05 (0.006)	-0.02 (0.006)	0.03 (0.006)	0.03 (0.006)	0.02 (0.006)
Well Age (months)	-0.013 (0.00007)	-0.026 (0.000146)	-0.037 (0.000259)	-0.048 (0.000415)	
Well Age ² (months)		0.00009 (0.000001)	0.00029 (0.000004)	0.00059 (0.00001)	
Well Age ³ (months)			-0.000001 (0.00000002)	-0.000004 (0.00000009)	
Well Age ⁴ (months)				0.00000001 (0.0000000003)	
N (Well-Months)	9,463,630	9,463,630	9,463,630	9,463,630	9,463,630
Number of Wells	157,578	157,578	157,578	157,578	157,578
Well Fixed Effects	✓	✓	✓	✓	✓
Cubic Spline					✓
R-Squared (Full Model)	0.662	0.669	0.671	0.672	0.672
R-Squared (Excluding Fixed Effects)	0.135	0.153	0.158	0.159	0.160

Notes: Clustered standard errors in parentheses. The first month of each well's production is dropped, as wells are typically operational for only a fraction of its first month. These results are robust to excluding natural gas prices from the specification.

Sources: Authors' calculations based on data from Drillinginfo and Bloomberg.

the additional supply from a 2015 baseline for different price levels and at different time horizons.

III.D.1 Simulation 1: Comparison of Unconventional and Conventional Price Response

In the first simulation, we calculate the time-varying supply response to a 10 percent price shock that raises oil prices from \$50 to \$55 per barrel. We assume baseline spud activity of 864 wells of each well type, which is equal to the average monthly spud count in our data in 2015. We use this baseline for both conventional and unconventional wells in order to compare them on a per-well basis, so that the choice of baseline does not affect the relative price response of each well type.²⁵ We calculate how many additional wells would be drilled over time due to the price shock using the preferred drilling response estimates from columns (1) and (4) of Table 2. These estimates resulted in cumulative responses of 1.19 for conventional and 1.63 for unconventional, respectively, suggesting an approximately 40 percent stronger drilling response for unconventional wells.

Given the simulated number of wells drilled in each month, we simulate how these wells would begin production over time according to the estimated survival functions from the preferred specifications in columns (1) and (3) of Table 3. This gives a simulated number of wells entering production in each month, from which we can calculate the total incremental oil production from wells of each type using the average decline curves portrayed in Figure 5.²⁶

The results of this simulation are shown in Figure 6. The top panel shows the percentage change in wells beginning production over time (i.e., wells that have been both drilled and

²⁵It is important to note that this assumption implies that this first simulation should not be interpreted as a forecast. Instead, it is meant to normalize the baseline to allow an apples-to-apples comparison of the price response of unconventional versus conventional wells on a per-well basis. For example, the actual conventional spud count that year was 326. Using that figure as a baseline for conventional would be a more accurate depiction of reality, but it would make it more difficult to compare and interpret the relative unconventional and conventional responses on a per-well basis.

²⁶Given the small elasticities in Table 4, we do not include any price response on this margin. Including such a response would make the unconventional/conventional difference somewhat larger. In addition, we model initial well productivity as fixed at 2015 levels. In principle, initial productivity itself may respond to a price change, but it is not clear whether this effect is positive or negative on net. On one hand, productivity may fall with higher prices, as less-productive wells become profitable. On the other hand, higher prices may encourage more effort by firms such as longer laterals and more intensive fracturing, causing productivity to rise. While the net of these effects is unclear *ex ante*, simple time series regressions of quarterly changes in log average initial production on changes in log oil prices yield small and statistically insignificant estimated elasticities on the order of -0.1 or smaller, depending on the specification. This suggests little evidence for a substantial net effect of prices on initial oil productivity, at least in the short term.

completed, incorporating the first two models).²⁷ The faster completion time of conventional wells mean that they ramp up slightly more quickly during the first couple of months. But unconventional wells quickly come online and overtake conventional ones due to the larger drilling response. After about two years, the rise in unconventional wells is about 40 percent larger due to the 40 percent larger estimated drilling response.

This 40 percent advantage for unconventional wells in terms of drilling responsiveness is magnified many times over by their much higher productivity. Despite taking somewhat longer to come online, unconventional wells produce about 9 times more oil than conventional wells immediately after coming online and 4.6 times larger cumulatively. This leads to the rapid rise in production shown in the bottom panel of Figure 6. While the incremental production continues to grow after that, it does so at a reduced rate as the steep decline curves reduce the contribution of the initial wells drilled. In the long-run, the total change in oil production in the bottom panel is about 6.4 times larger for unconventional wells (an increase of about 13 percent of baseline production, compared to about 2 percent for conventional, given a 10 percent price increase). This 6.4-fold advantage reflects the combination of unconventional wells' 40 percent stronger drilling response and their 4.6-fold long-run higher productivity, mentioned in section II (that is, $6.4 = 1.4 \times 4.6$).²⁸ While the stronger drilling response plays some role, most of the larger price response owes to unconventional wells' higher productivity.

This simulation directly compares unconventional and conventional wells on a per-well basis, revealing the sources of unconventional wells' larger responsiveness. However, it involves making normalizing assumptions that do not fully reflect changes to the composition

²⁷This is a percent of *total* wells, unconventional and conventional. With a baseline of 864 wells of each type, this baseline is $864 + 864 = 1,728$. Since the denominators include both types of wells in equal measure, the total response is half as large would be obtained by a simple calculation of the cumulative elasticity (1.6 and 1.2) times the price change (10 percent). To compute the percentage change in drilling relative to each type's *own* baseline (864), one simply needs to multiply the values in the top panel of Figure 6 by 2.

²⁸A slightly simplified version of this simulation shows why the spud and productivity advantages combine multiplicatively. By construction, the long-run oil supply response equals the productivity per-well times the long-run change in wells drilled, denoted $r_j := \tilde{q}_{oil,j} \Delta w_j$ for $j \in \{u, c\}$. Given our log specification in equation (1), the long-run change in wells drilled is approximately equal to the product of the percentage price change ($\Delta \hat{p} := \frac{\Delta p_{oil}}{p_{oil}}$), the long-run drilling response ($\tilde{\beta}_j := \sum_{l=0}^L \beta_{1,j,l}$) and the baseline number of wells drilled (w_j^{base}). This gives $\Delta w_j = \Delta \hat{p} \times \tilde{\beta}_j \times w_j^{base}$. Plugging this into the expression for the oil response gives $r_j = \tilde{q}_{oil,j} \times \Delta \hat{p} \times \tilde{\beta}_j \times w_j^{base}$. The relative advantage of unconventional over conventional is the ratio of the unconventional and conventional production responses, r_u and r_c , which can be written as

$$\frac{r_u}{r_c} = \frac{\tilde{q}_{oil,u}}{\tilde{q}_{oil,c}} \times \frac{\tilde{\beta}_u}{\tilde{\beta}_c} \times \frac{w_u^{base}}{w_c^{base}}. \quad (3)$$

The overall production response ratio is thus written as the product of three ratios: the long-run productivity ratio (4.6), the cumulative price response ratio (1.4), and the ratio of baseline drilling activity (which were assumed to be the same for this simulation to compare the well types on an equal footing, making this ratio 1).

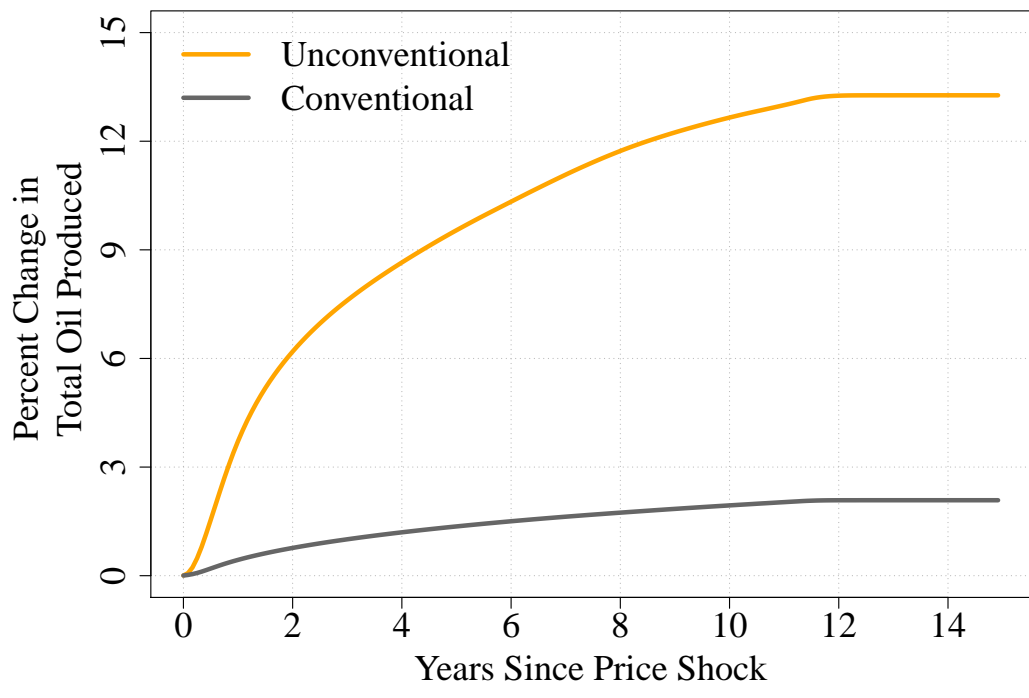



Figure 6: Percent Change in Oil Wells Beginning Production and Oil Produced from Those Wells, Following a 10 Percent Price Shock

Sources: Authors' calculations based on data from Drillinginfo and Bloomberg.

of U.S. oil supply and the oil price environment over time. Importantly, it does not capture the much larger number of unconventional wells drilled in recent years, which affects the overall price responsiveness. Our second simulation accounts for these factors.

III.D.2 Simulation 2: Comparison of Overall Price Response Over Time

In the second simulation, we consider how the changes in the composition of U.S. oil supply in the past decade have altered its overall responsiveness and the degree to which it has changed the role of the United States as a “swing producer.” We do this by running the previous simulation eleven times, each time starting in each year between 2005 and 2015. For each starting year, we simulate a \$10 price shock (which is typically larger than the 10 percent price shock above), and we calibrate each simulation to the year’s observed oil prices, average nationwide oil drilling activity, and changing initial productivity of both conventional and unconventional oil wells. By doing so, the simulation effectively captures the shift towards unconventional drilling in recent years. This also captures the fact that baseline drilling activity will be different in low-price environments compared to high-price ones. The baseline drilling activity values are computed separately for unconventional and conventional wells to account for the changing composition of oil supply during this period.²⁹

The results of these simulations are shown in Figure 7.³⁰ The lines are shaded in a gray-to-orange hue in proportion to the share of baseline drilling activity that was dominated by conventional (gray) or unconventional (orange) in that year.

The top panel shows the change in new wells entering production over time, by simulation year. In the years pre-dating the shale oil boom (2005-2009), the response to a \$10 price increase after one year is approximately 115 incremental new wells each month, most of which are conventional (as indicated by the gray shading) due to the fact that shale techniques had not yet been widely applied to oil wells. Turning to the oil produced by these wells, shown in the bottom panel, these wells would eventually increase total oil supply by about 100,000 barrels per day after about five years.

After the shale oil boom begins in 2010, the simulated ramp-up in drilling is much larger (on the order of 175 to 300 incremental wells per month). Further, this drilling response is increasingly focused on more productive shale wells (indicated by a deeper orange color),

²⁹We calculated nationwide spud activity in each year based on a supplementary dataset downloaded from Drillinginfo. For all parameters other than baseline spud counts, we use the means of the variables in our sample for the relevant time period (namely, prices and initial production).

³⁰Note that the horizontal axis in the bottom panel of this simulation is truncated at 5 years (compared to the corresponding panel in Figure 6, which extended to 15 years, since that is when the system reaches equilibrium). We truncate the axis here to focus on short-to-medium-run effects, which are key to the question of whether the United States is a new “swing producer”.

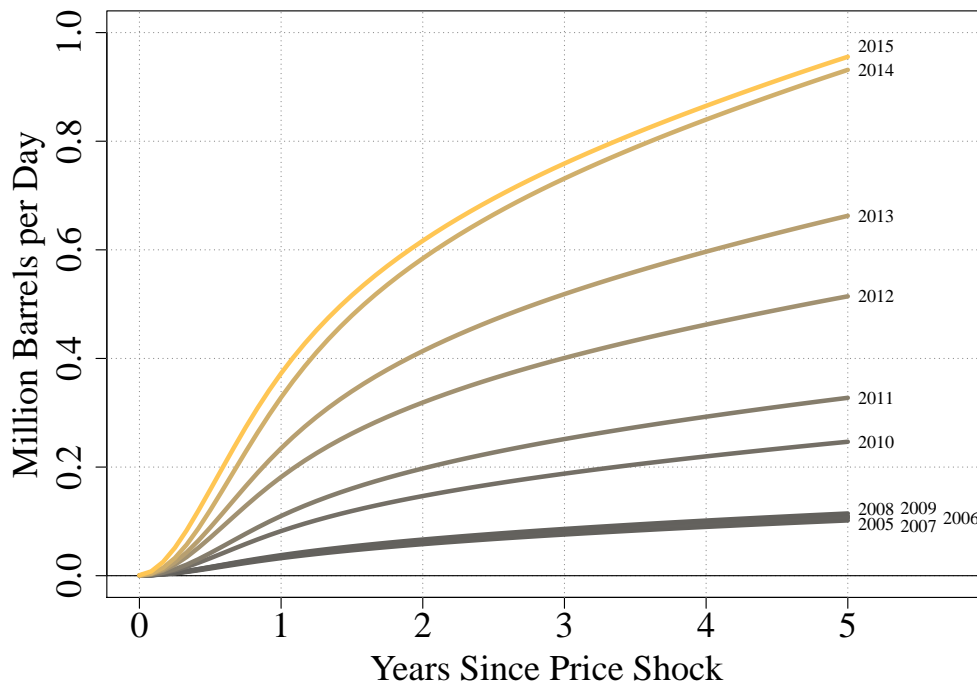
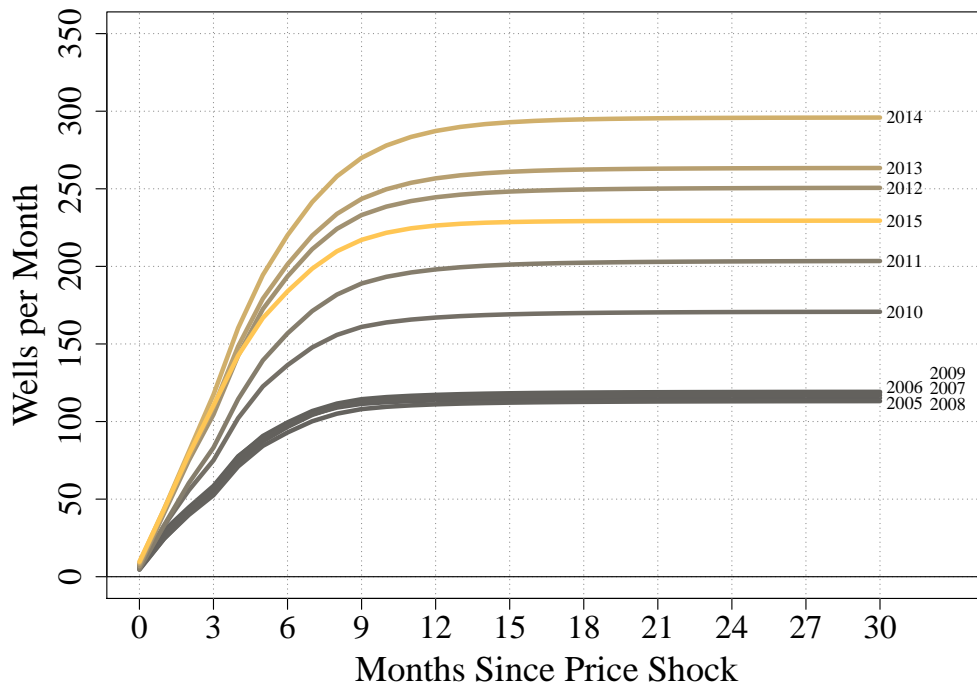


Figure 7: Change in Oil Wells Beginning Production and Oil Produced from Those Wells, Following a \$10 Price Shock, by Year of Price Shock

Sources: Authors' calculations based on data from Drillinginfo and Bloomberg.

leading to a much larger oil supply response. By 2015, the simulated supply response to the \$10 price shock is on the order of 400,000 barrels per day after 1 year, 600,000 barrels after 2 years, and nearly 1 million barrels after 5 years.

This larger production response occurs despite a smaller drilling response in the 2015 simulation (top panel), particularly as compared to 2014. This difference is not due to an increase in well productivity by either unconventional or conventional wells; indeed, Figure A.1 shows no substantial increase in productivity that year for either type of well. Rather, the dichotomy between the production and drilling response is due to a refocusing of the industry on higher-production unconventional drilling, and away from conventional drilling.

The production response in the 2015 simulation represents an 8.7-fold larger supply response compared to the 2005 simulation (960,000 versus 110,000). This differs from the 6.4-fold larger price response from the first simulation presented in Figure 6 because it also takes into account the changing the per-well productivity, price environment, composition of U.S. oil supply (recall the first simulation normalizes this factor to conduct an apples-to-apples, per-well comparison).

Nonetheless, over shorter time frames (3-6 months), the increases in supply are restricted due to the time needed drilling activity to ramp up and for drilled wells to be fractured, completed, and begin production. Even in the 2015 simulation, the incremental supply is modest (50,000 barrels per day or less) in the first three months following the price shock. Six months after the price shock, production ramps up to approximately 160,000 barrels per day. This is substantial, particularly compared to the negligible response over a six-month period in the 2005 simulation; at that time an increase of more than 100,000 barrels per day would take several years, instead of six months. Nonetheless, even in 2015 it takes six months for this supply to appear. This timing is inconsistent with the traditional notion of a “swing producer”, which is a supplier able to quickly respond to emergency supply disruptions within 30 to 90 days.

III.D.3 Simulation 3: Price Response by Price Change and Time Horizon

The third set of simulations are illustrated in Figure 8. In these simulations, we fix the baseline to 2015 levels using the same approach as in the previous section, but we vary the price change. The figure shows the production response at price levels ranging from \$50 (no change) to \$80 at different time horizons (6 months, 1 year, 2 years, and 5 years). This is meant to represent our model’s indication for how much supply could be brought online at each time horizon, given an increase in prices from a baseline value of \$50 per barrel.

The lines show the production response at different price levels. For example, the top of the line on the far left indicates that if prices rise to \$80, production would be higher (relative

to a \$50 counterfactual) by about 500,000 barrels per day after six months, due to the time lags between changes in drilling activity and changes in barrels produced. After a year at \$80 oil however, production would be higher by 1.2 million barrels per day (MMBPD), 2 MMPBD after 2 years, and 3 MMBPD after 5 years.

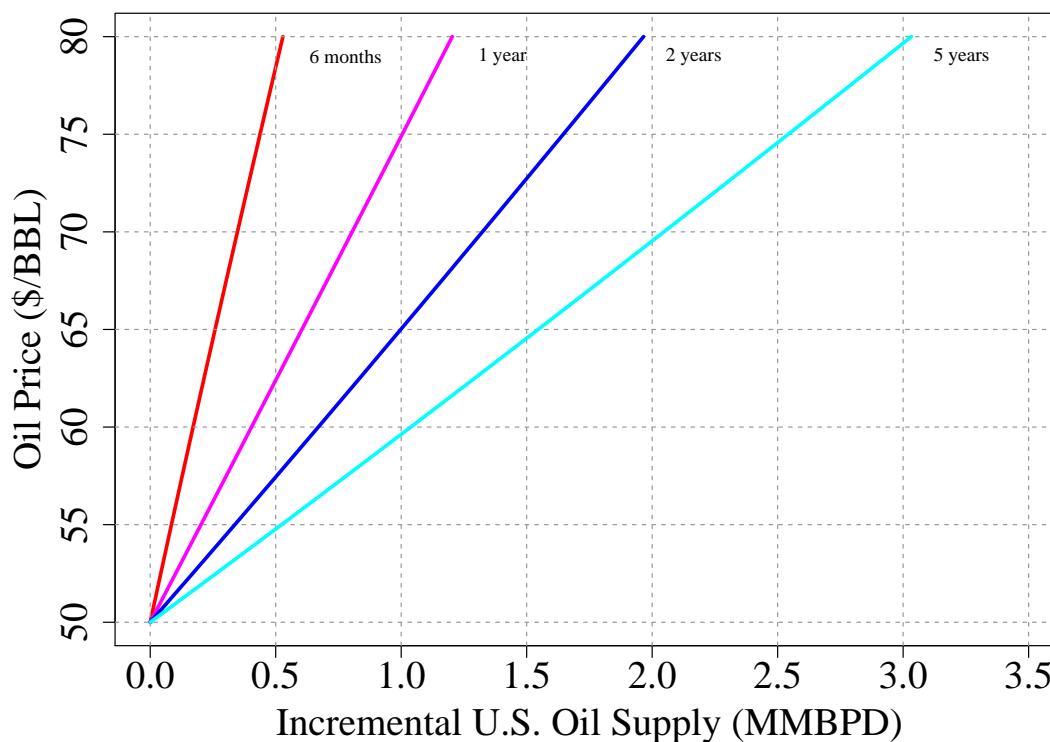


Figure 8: Price versus Incremental Oil Supply Following a Price Increase Relative to \$50, by Time Horizon

Sources: Authors' calculations based on data from Drillinginfo and Bloomberg.

IV Conclusion

We empirically assess the differences in price responsiveness for oil wells using a detailed dataset on 164,000 oil wells during 2000-2015 in five major oil-producing states of Texas, North Dakota, California, Oklahoma, and Colorado. We estimate the price responsiveness at three key stages of production: drilling, spud-to-production time, and production from existing wells. We find that the important margin for the price response is drilling activity, estimating a somewhat higher drilling price responsiveness of 1.6 per 1 percent change in oil prices for unconventional wells compared to 1.2 for conventional wells.

We find that unconventional oil wells are more price responsive mostly due to their much higher productivity (approximately 9 times larger immediately after beginning production, and 4.6 times larger cumulatively in the long run). We conduct simulations to combine our different models to show how the rise of unconventional drilling has affected the aggregate supply response. The slightly larger estimated drilling advantage combined with the larger oil produced per well leads to an estimated 6-fold larger price response from unconventional oil wells on a per-well basis. Further accounting for the sharp rise in unconventional drilling (compared to conventional drilling) and changing market environment in price environment, drilling composition, and well productivity makes this difference even larger, implying a 9-fold larger price response compared to the pre-shale era.

Our simulations suggest that if oil prices were to rise from \$50 to \$80 per barrel, U.S. production could ramp up production by 0.5 million barrels per day in 6 months, 1.2 million in 1 year, 2 million in 2 years, and 3 million in 5 years. These represent significant increases in context of the global market, suggesting a significantly larger role for U.S. incremental supply than before the shale revolution. However, the response still takes more time to arise than is typically considered for a “swing producer”, referring to a supplier able to increase production quickly, within 30 to 90 days.

References

- Anderson, Soren T., Ryan Kellogg, and Stephen W. Salant.** 2017. "Hotelling Under Pressure." *Journal of Political Economy*.
- Barsky, Robert B, and Lutz Kilian.** 2001. "Do We Really Know That Oil Caused the Great Stagflation? A Monetary Alternative." *NBER Macroeconomics Annual*, 16: 137–183.
- Baumeister, Christiane, and James 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*.
- Baumeister, Christiane, and Lutz Kilian.** 2012. "Real-Time Forecasts of the Real Price of Oil." *Journal of Business & Economic Statistics*, 30(2): 326–336.
- Baumeister, Christiane, and Lutz Kilian.** 2016a. "Forty Years of Oil Price Fluctuations: Why the Price of Oil May Still Surprise Us." *The Journal of Economic Perspectives*, 30(1): 139–160.
- Baumeister, Christiane, and Lutz Kilian.** 2016b. "A General Approach to Recovering Market Expectations from Futures Prices with an Application to Crude Oil." *manuscript*.
- Coglianese, John, Lucas W Davis, Lutz Kilian, and James H Stock.** 2017. "Anticipation, Tax Avoidance, and the Price Elasticity of Gasoline Demand." *Journal of Applied Econometrics*.
- Collins, Gabriel, and Kenneth B. Medlock.** 2017. "Assessing Shale Producers Ability to Scale-up Activity." Rice University's Baker Institute for Public Policy, Houston, Texas Issue brief no. 01.17.17.
- Covert, Thomas.** 2015. "Experiential and Social Learning in Firms: The Case of Hydraulic Fracturing in the Bakken Shale." *Available at SSRN 2481321*.
- Dahl, Carol, and Mine Yücel.** 1991. "Testing alternative hypotheses of oil producer behavior." *Energy Journal*, 117–138.
- Davis, Lucas W, and Lutz Kilian.** 2011. "Estimating the effect of a gasoline tax on carbon emissions." *Journal of Applied Econometrics*, 26(7): 1187–1214.
- Griffin, James M.** 1985. "OPEC behavior: a test of alternative hypotheses." *American Economic Review*, 954–963.
- Güntner, Jochen HF.** 2014. "How do oil producers respond to oil demand shocks?" *Energy Economics*, 44: 1–13.
- Hamilton, James D.** 2009. "Understanding Crude Oil Prices." *Energy Journal*, 30(2): 179–206.
- Hogan, William W.** 1989. *World oil price projections: a sensitivity analysis*. Harvard University, Energy and Environmental Policy Center, John F. Kennedy School of Government.

- Jones, Clifton T.** 1990. "OPEC behaviour under falling prices: implications for cartel stability." *Energy Journal*, 117–129.
- Joskow, Paul L.** 2013. "Natural Gas: From Shortages to Abundance in the United States." *American Economic Review Papers and Proceedings*, 103(3): 338–43.
- Kahle, David, and Hadley Wickham.** 2013. "ggmap: Spatial Visualization with ggplot2." *The R Journal*, 5(1): 144–161.
- Kilian, Lutz.** 2009. "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market." *American Economic Review*, 99(3): 1053–69.
- Kilian, Lutz.** 2016. "The Impact of the Shale Oil Revolution on U.S. Oil and Gasoline Prices." *Review of Environmental Economics and Policy*, 10(2): 185–205.
- Kilian, Lutz.** 2017a. "How the Tight Oil Boom Has Changed Oil and Gasoline Markets." *CEPR Discussion Paper No. DP11876*.
- Kilian, Lutz.** 2017b. "The Impact of the Fracking Boom on Arab Oil Producers." *CESifo Working Paper Series No. 5751*.
- Newell, Richard G., Brian C. Prest, and Ashley Vissing.** 2016. "Trophy Hunting vs. Manufacturing Energy: The Price-Responsiveness of Shale Gas." National Bureau of Economic Research Working Paper 22532.
- Ramcharan, Harri.** 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, James L.** 2017. "Estimating the Future Supply of Shale Oil: A Bakken Case Study."

A Appendix

A.1 Drilling IV First Stage

Table 5 contains the first stage regressions for the 2SLS regressions in Table 2.

Table 5: First Stage Regressions Drilling Estimation

	Conventional Regressions		Unconventional Regressions	
	<i>Dependent variable:</i>			
	Oil Revenues	Oil Prices	Oil Revenues	Oil Prices
	(1)	(2)	(3)	(4)
$\Delta\text{Log}(\text{CRB Index})$	1.58 (0.21)	1.40 (0.18)	1.10 (0.25)	1.43 (0.18)
$\Delta\text{Log}(\text{Gas Revenues or Prices})$	0.05 (0.09)	0.15 (0.09)	0.20 (0.09)	0.12 (0.12)
$\Delta\text{Log}(\text{Gas Revenues or Prices}), 1 \text{ Lag}$	0.09 (0.06)	0.14 (0.08)	0.29 (0.06)	0.29 (0.11)
$\Delta\text{Log}(\text{Gas Revenues or Prices}), 2 \text{ Lags}$	0.05 (0.08)	0.02 (0.06)	0.17 (0.12)	0.09 (0.10)
$\Delta\text{Log}(\text{Gas Revenues or Prices}), 3 \text{ Lags}$	0.04 (0.05)	-0.01 (0.07)	0.04 (0.08)	-0.01 (0.11)
$\Delta\text{Log}(\text{Oil Revenues or Prices}), 1 \text{ Lag}$	0.14 (0.08)	0.07 (0.09)	0.01 (0.09)	0.04 (0.10)
$\Delta\text{Log}(\text{Oil Revenues or Prices}), 2 \text{ Lags}$	0.05 (0.08)	-0.01 (0.08)	0.07 (0.11)	-0.07 (0.07)
$\Delta\text{Log}(\text{Oil Revenues or Prices}), 3 \text{ Lags}$	-0.01 (0.11)	0.08 (0.10)	0.01 (0.09)	0.01 (0.10)
2nd Quarterly Lag	0.14 (0.04)	0.13 (0.04)	0.21 (0.08)	0.16 (0.05)
3rd Quarterly Lag	0.06 (0.05)	0.10 (0.05)	0.19 (0.08)	0.10 (0.07)
4th Quarterly Lag	0.01 (0.05)	0.05 (0.04)	0.05 (0.05)	0.03 (0.04)
Constant	-0.06 (0.03)	-0.07 (0.03)	-0.11 (0.05)	-0.08 (0.04)
Observations	63	63	43	43
R ²	0.52	0.56	0.64	0.66
Adjusted R ²	0.42	0.46	0.51	0.54
F Statistic (Instrument Only)	54.7	58.1	19.1	63.1

Note: Dependent variable is the $\Delta\text{Log}(\text{Oil Revenues})$ or $\Delta\text{Log}(\text{Oil Prices})$. Columns (1) and (2) differ from (3) and (4) because of the different sample period (2000-2015 versus 2005-2015) and revenue variables (conventional versus unconventional).

Sources: Authors' calculations based on data from Drillinginfo and Bloomberg.

A.2 Oil Well Productivity Over Time

Figure A.1 shows the trends in productivity for both unconventional and conventional wells. This represents a two-quarter moving average of the first full month of oil production. These are the same variables that are multiplied by oil prices to construct our variables representing oil revenues.

The productivity of conventional wells has been mostly flat since 2000, at just under 50 barrels per day.

For unconventional wells, initial productivity has been rising steadily, from 150 barrels per day in 2005 to nearly 400 barrels in 2015. The roughly linear nature of the rise suggests that the rise is attributable to exogenous technological change, not by swings in oil prices altering drilling behavior.

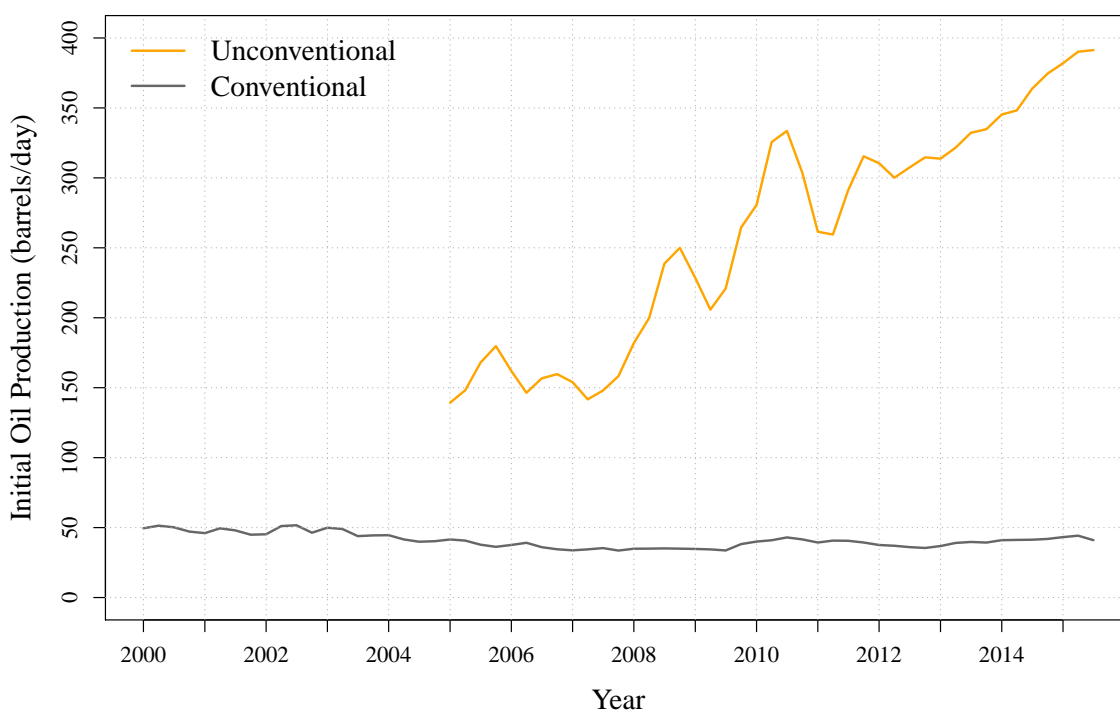


Figure A.1: Average Oil Production Per Well During the First Full Month, 2000-2015, Quarterly

Sources: Authors' calculations based on data from Drillinginfo.

A.3 Oil Decline Curves as a Percentage of Initial Production

Figure A.2 shows the average decline curves from Figure 5 scaled as a percent of initial production. This allows one to compare the percentage declines across unconventional and conventional wells. This reveals that unconventional oil wells decline much faster than conventional ones. For example, after 12 months, unconventional wells have declined by about 70 percent, compared to only about 50 percent for conventional wells.³¹ This highlights how steep decline curves are a distinguishing feature of shale oil wells, as has been commonly discussed.




Figure A.2: Mean and Median Profiles of Monthly Oil Production from Oil Wells, as a Percent of Initial Production

Sources: Authors' calculations based on data from Drillinginfo.

³¹This may appear to be inconsistent with the decline rates shown in Table 1. This apparent inconsistency is due to the fact the decline rates are calculated differently. For example, the 12-month decline rates in Table 1 are computed as the average across well-specific 12-month decline rates as calculated by Drillinginfo. This calculation requires dropping wells that have not yet produced for 12 months. In addition, Drillinginfo's decline rate calculations often involve some rounding error, which can be significant. In contrast, this figure uses more detailed time series information from all wells, however long they have been producing, which better represents the true decline curve.

A.4 Distribution of Initial Production, by Well Type

Figure A.3 shows for each well type the distribution of initial production (daily oil production during the first full month of production). The top panel shows the distributions for all wells in our data, whereas the bottom panel excludes wells that produced zero oil in their first full month. (Many such wells often eventually produced oil, but they were typically less productive on average.) The distributions are truncated at the 99th percentile for readability.

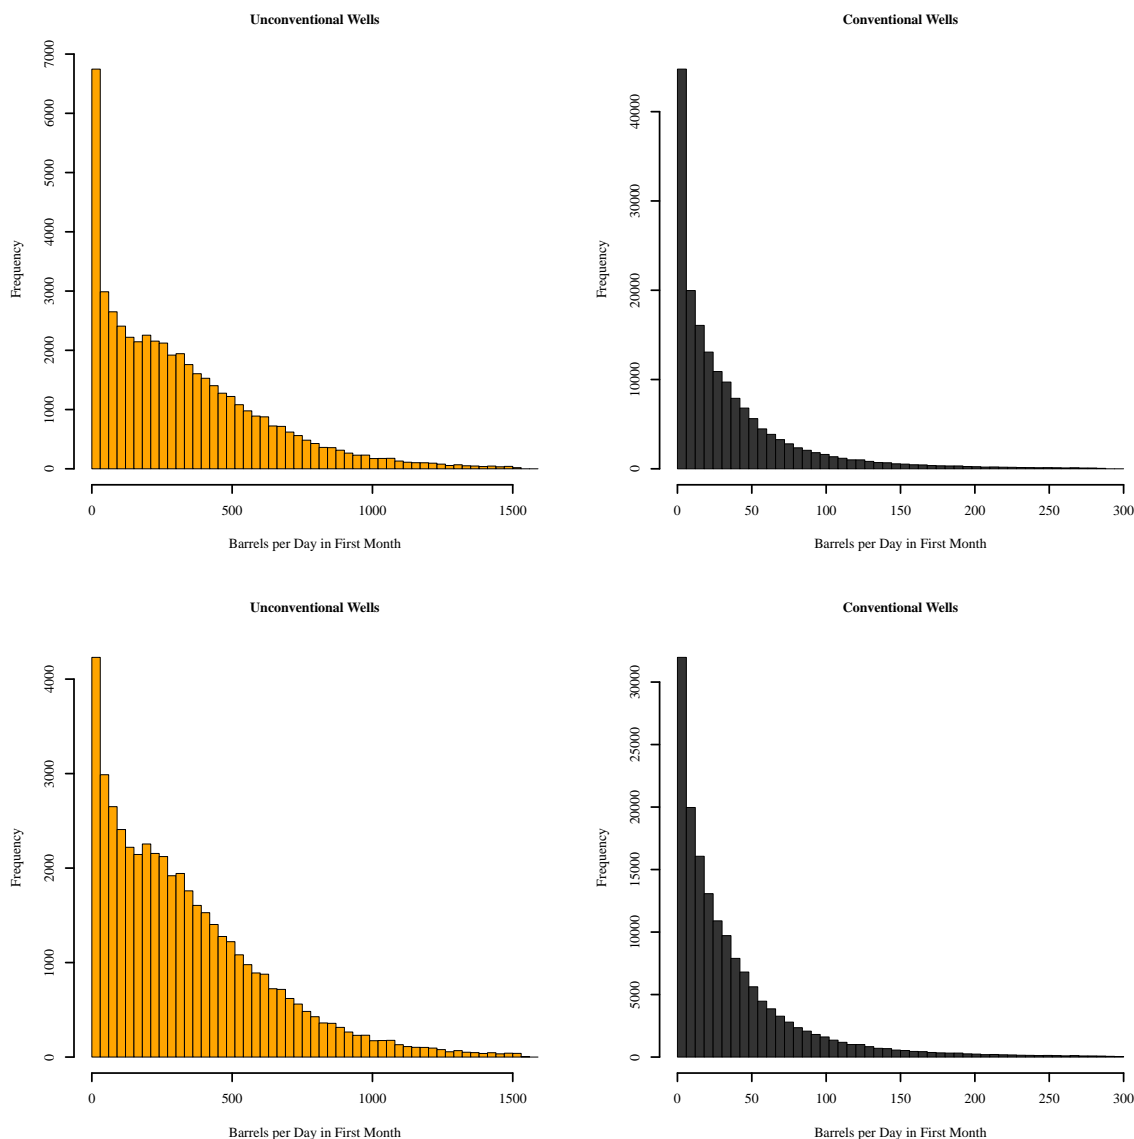


Figure A.3: Distributions of Initial Production (IP) by Well Type, Including (Top) and Excluding (Bottom) Wells with Zero IP

Sources: Authors' calculations based on data from Drillinginfo.