

CAMP Working Paper Series
No 5/2021

The Price Responsiveness of Shale Producers: Evidence From Micro Data

Knut Are Aastveit, Hilde C. Bjørnland and
Thomas S. Gundersen



© Authors 2021 This paper can be downloaded without charge from the CAMP website bi.no/camp



The Price Responsiveness of Shale Producers: Evidence From Micro Data*

Knut Are Aastveit[†] Hilde C. Bjørnland[‡] Thomas S. Gundersen[§]

September 11, 2021

Shale oil producers respond positively and significantly to favourable oil price signals. This finding is established using a novel proprietary data set consisting of more than 200,000 shale wells across ten U.S. states spanning almost two decades. We document large heterogeneity in the estimated responses across the various shale wells, suggesting that aggregation bias is an important issue for this kind of analysis. We find responses to be stronger for the largest oil producing firms, among wells that are spaced further apart and in regions where the density of shale wells is higher. The response also depend on the level of production. Our empirical results calls for new models that can account for a growing share of shale oil in the U.S., the inherent flexibility of shale extraction technology in production and the role of shale oil in transmitting oil price shocks to the U.S. economy.

JEL-codes: C23, Q41, Q43

Keywords: Oil price, Shale oil supply, Well-level panel data

*The authors would like to thank Christiane Baumeister, Jamie Cross, Jon H. Fiva, Ivan Petrella, Leif Anders Thorsrud, seminar participants at the November 2019 Oslo Macro Group meeting at BI Norwegian Business School, as well as participants at the 2019 MRG workshop at the University of Sydney for valuable comments and suggestions. This paper is part of the research activities at the Centre for Applied Macroeconomics and Commodity Prices (CAMP) at the BI Norwegian Business School. The usual disclaimers apply. The views expressed in this paper are those of the authors and do not necessarily reflect those of Norges Bank.

[†]Norges Bank and Centre for Applied Macroeconomics and Commodity Prices (CAMP), BI Norwegian Business School. Email: knut-are.aastveit@norges-bank.no.

[‡]Centre for Applied Macroeconomics and Commodity Prices (CAMP), BI Norwegian Business School and Norges Bank. Email: hilde.c.bjornland@bi.no.

[§]Corresponding author: Centre for Applied Macroeconomics and Commodity Prices (CAMP), BI Norwegian Business School. Email: thomas.s.gundersen@bi.no.

1 Introduction

Theoretical models and empirical analysis of conventional oil extraction find no response in production to oil price changes. Instead, they find that conventional oil wells are constrained by reservoir pressure and gradually decline, see for instance Anderson, Kellogg, and Salant (2018) for a general model applied to Texas, or Pesaran (1990), Dahl and Yucel (1991), Ramcharran (2002) and Smith (2009) for empirical studies of conventional oil production across many U.S. states. Over the past decade, however, the oil market has undergone a significant transformation caused by an unexpectedly sharp increase in U.S. crude oil and natural gas production from unconventional (shale) wells.¹ This massive production surge of shale oil and gas is made possible by the development of hydraulic fracturing (so-called "fracking") and horizontal drilling technologies, making the U.S. the world's largest oil and natural gas producer. A key feature of fracking is that it allows for a more flexible production process compared to conventional oil production, as wells can be re-fractured over time. This implies that oil companies can be forward looking, reducing the extraction rate when market conditions are poor, or resuming extraction when conditions improve, see Bornstein, Krusell, and Rebelo (2021) for a structural model of the global oil market that incorporates fracking. Hence, we should expect fracking firms to be more responsive to changes in oil prices than conventional oil producing firms.

While empirical studies of conventional oil production are ample, there is a scarcity of empirical studies analyzing production from shale wells, mostly due to lack of high frequent data at the well level. In this paper, we aim to fill this gap by studying the price responsiveness of U.S. oil producers using a novel proprietary dataset compiled by Rystad Energy. The data set contains monthly information on crude oil production and other characteristics for more than 200,000 horizontally drilled shale wells. The data set covers production from all reported shale oil wells in the 10 largest U.S. oil producing states for the period 2005:M01–2017:M12. In addition to this, we also have access to data on more than 150,000 conventionally drilled wells in Texas for the same time period. To our knowledge, this makes our study the most detailed and comprehensive study on the behaviour of U.S. shale oil producers to date.

We start by analyzing production at the well level, using the detailed cross-section of our micro data. To our knowledge, only two studies have analyzed the price-responsiveness of U.S. shale oil producers using high frequent data at the well level, and they reach different conclusions: Bjørnland, Nordvik, and Rohrer (2021) examine oil producers in North Dakota and finds a positive price response for well-completion and production

¹Shale oil production involves pumping a mixture of liquids and sand at high pressure into shale rock formations with low permeability to release oil and gas trapped in small pockets. This is combined with the ability to drill horizontally through shale layers over long distances.

from existing shale wells, while Newell and Prest (2019) analyze five major oil producing states including North Dakota, but finds no response for shale production to price signals, only drilling and completions respond positively. Apart from the differences in datasets, the two studies employ notably different modelling frameworks.² Our baseline model is specified so as to capture features of both models. In particular, the model is specified in log-level as Newell and Prest (2019), and contains well age and well fixed effects.³ Furthermore, we estimate the response of U.S. oil producers to both the spot price of oil and the oil spot-futures spread as Bjørnland, Nordvik, and Rohrer (2021), and control for the response to natural gas prices as Newell and Prest (2019).

We estimate the model using level information of the ten oil producing states in our sample taken together, before analyzing if there is heterogeneity across states and shale plays by estimating our model separately for each. Doing so, we document large differentials in price responsiveness between shale wells and conventional wells. Consistent with earlier findings, such as Anderson, Kellogg, and Salant (2018), we find no response of conventional wells in Texas. In contrast, we find that shale oil producers respond significantly to price signals. While the response to the current spot price is small, shale producers respond strongly to movements in the spot-futures spread, by increasing their production on impact when the spot-futures spread increases. Furthermore, we show that for the most part, there is a large positive and statistically significant response irrespective of state. The exceptions are California, New Mexico and Wyoming. The lack of response from these can be understood through geographical isolation or lack of sufficient pipeline infrastructure. In the shale plays, the responsiveness tends to be stronger than in the state the play is located, especially in the case of Texas. Our results confirm the findings for North Dakota by Bjørnland, Nordvik, and Rohrer (2021), using different models and sample.

Since shale oil extraction technology introduces new decision margins with well completion and refracturing of existing wells, we explore a series of additional hypothesis. First, we explore if the price responsiveness depend on the production level. Since shale wells are characterised by front-loaded production profiles, the level of production is very high in months immediately following the initiation of production or the months following a refracturing event. Using quantile regressions, we find supporting evidence that it is when well output is in the upper right tail of the distribution that they are the most

²Bjørnland, Nordvik, and Rohrer (2021) base the analysis on growth rates and the use of directional movements in relative oil futures prices to estimate a supply elasticity. Newell and Prest (2019) on the other hand, base the analysis on a mixture of futures prices and spot prices for oil and natural gas in levels of the variables. Specifically, futures prices are used in the analysis of drilling intensity while spot prices are employed for the estimation of a supply elasticity for already producing wells.

³Our results are robust to estimating the model in log-differences instead of log-levels.

price responsive. This corresponds to completion and refracturing events. Second, we examine if responses depend on well ownership. Identifying well operators in our dataset, we find that well price-responsiveness also tends to be stronger if the well is owned and operated by one of the large, and most likely most professional, firms. Finally, for technical reasons, wells that are spaced closely together may interfere with each other because they tap into the same reservoir. An appropriate well spacing is therefore desirable. This issue may be exacerbated by fracking because it is not an easily controlled process and fractures can be made into adjacent wells. We thus investigate whether well spacing can be a determinant for how strong the price response is. Indeed, we find that wells spaced farther apart exhibit a large price-response compared to those that are spaced less than 600 feet (approx. 183 metres) apart.

The use of micro data to analyze price responses of aggregate oil production has some advantages. In particular, by constructing a panel dataset based on rich well level information, we can eliminate any potential aggregation bias over well production rates when estimating the empirical model. Aggregating over all individual wells across the U.S. is equivalent to imposing identical parameter values for all producing wells regardless of well or firm characteristics. In fact, when we aggregate production across individual wells in our panel and estimate price responses we find that aggregate output is price-inelastic for both shale and conventional oil. Hence, doing so we would misleadingly infer that shale producers have no response on impact to oil price signals. Furthermore, the use of panel data enables us to explore the cross-sectional variation in, for instance, well type, well operator, location, or other characteristics of interest, and we can investigate the potential heterogeneity in producer behavior. Lastly, having a large cross section in a panel is beneficial for statistical inference when analyzing a relatively short time period as we do here.

Our results have important macroeconomic implications. First, as shale producers are forward looking, we may expect to see a stabilizing effect on oil prices as shale producers grow in size and importance, see Bornstein, Krusell, and Rebelo (2021). For instance, a persistent increase in the oil price (due to say increased global demand) will now make it profitable to expand shale production and take advantage of the high oil prices, thereby bringing the oil price effectively down again. Second, policymakers should take into account that shale and conventional producers adjust differently to price-sensitive news and policies. In particular, the increased flexibility among shale producers could have implications for whether supply side policies are effective. For instance, it has been argued that cuts in oil production can have positive climate effects (i.e., reducing CO₂ emissions) if it is not replaced by increased oil production elsewhere, see Fæhn, Hagem, Lindholt, Mæland, and Rosendahl (2017). However, as shale firms are forward looking and can

respond quickly, a cut in oil production in one location (which will shift the supply curve leftwards and increase oil prices immediately, all else equal) can be replaced by shale oil production elsewhere, at least in the short term. Hence, if supply side policies shall be effective, they need to be accompanied by some form of commitment/agreement among both the shale and conventional producers, see Asheim, Fæhn, Nyborg, Greaker, Hagem, Harstad, Hoel, Lund, and Rosendahl (2019). Third, the growing share of shale oil in the U.S. has implications for how investment, wages, employment and other macroeconomic variables may respond to oil price changes, see for instance Allcott and Keniston (2017) and Bjørnland and Zhulanova (2018) for some early studies. Related to this, our results can also serve to reconcile some of the opposing conclusions in the literature when it comes to how one should analyse the role of oil in the macroeconomy. In particular, oil price-macro models have often assumed aggregate oil production to be price inelastic in the short run when identifying oil market shocks, see for instance Kilian (2009) and Kilian and Murphy (2012). 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 strategies for oil market macro models that relax the assumption of a zero short-run oil supply elasticity, for instance Baumeister and Hamilton (2019) and Caldara, Cavallo, and Iacoviello (2019). To conclude, our empirical results call for new models that can account for a growing share of shale oil in the U.S., the inherent flexibility of shale extraction technology in production, and the role of shale oil in transmitting oil price shocks to the U.S. economy.

The rest of the paper is organized as follows. Section 2 provides background and characteristics of shale oil production and details about the micro data. In Section 3 we present the model and the empirical results. Specifically, Subsection 3.2 presents our baseline results for shale and conventional wells. In subsection 3.3, we present our quantile regression results and discussion of completion and refracturing events. Subsection 3.4 explores different dimensions in the cross-section of our data, namely with respect to well ownership and well spacing. The final Subsection 3.5 contains our results for the different states and shale plays and we then conclude in Section 4.

2 Background and data

2.1 Background

The renewed interest in oil supply responsiveness comes in part from the emergence of horizontally drilled shale wells in the United States. Contrary to the conventionally drilled vertical wells, shale wells can economically tap into vast shale rock formations that are

well known. These formations are of such low permeability that the hydrocarbons are trapped within tiny pockets from which they cannot escape without external stimulation. Conventional wells on the other hand, have a high degree of permeability, meaning that a well will naturally begin to flow if the pressure differential is made sufficient. The emergence of hydraulic fracturing (*fracking*) technology has made it possible to increase permeability of shale rocks. After a shale well is drilled, the owner contracts with a fracking crew who pumps the well full of water, chemicals and proppants at high pressures to create or expand rock fissures so that liquids (or natural gas) can flow. The role of proppants is to make sure that the created fissures stay open after the fracking is completed. This process is sometimes referred to as *well completion* because it completes the development phase of a well. Importantly, owners of shale wells have the option to postpone completion of a well allowing them to better time the decision to produce given overall market conditions. This option is lucrative because the average shale well outputs hundreds percent more crude oil during the first production months compared to conventional oil wells giving operators major incentives to optimise the timing of well completion. This increase in productivity compared to conventional wells is achieved by combining hydraulic fracturing with horizontal drilling technology. By turning the drill bit horizontally at the desired well depth, optimally in the shale rock layer, the well can tap into a larger surface area and thus increase yield. Furthermore, as the flow rate of a producing well diminishes, the producers have the option to re-stimulate a well to increase the expected ultimate recovery. This difference in production technologies leaves shale well operators much closer to the original Hotelling model behaviour with well completion being the decision variable. One should therefore expect shale well operators to be more sensitive to oil price changes than what is the case for conventional wells.

2.2 Data

The dataset that makes up the foundation of our analysis is a well-level panel at the monthly frequency covering all reported onshore oil wells producing in the ten major oil-producing contiguous U.S. states. These states are California, Colorado, Kansas, New Mexico, North Dakota, Montana, Oklahoma, Texas, Utah and Wyoming. The data is provided by Rystad Energy and cover more than 200,000 unique horizontally drilled wells.⁴ Figure 1 provides a plot of the geographical distribution of the wells in our dataset. We also have access to data on more than 150,000 conventionally drilled wells in Texas during the same time period. Our sample period runs from 2005M01–2017M12.

In addition to information about the number of barrels of crude oil produced in a

⁴Rystad Energy is an independent energy market intelligence firm headquartered in Oslo, Norway.

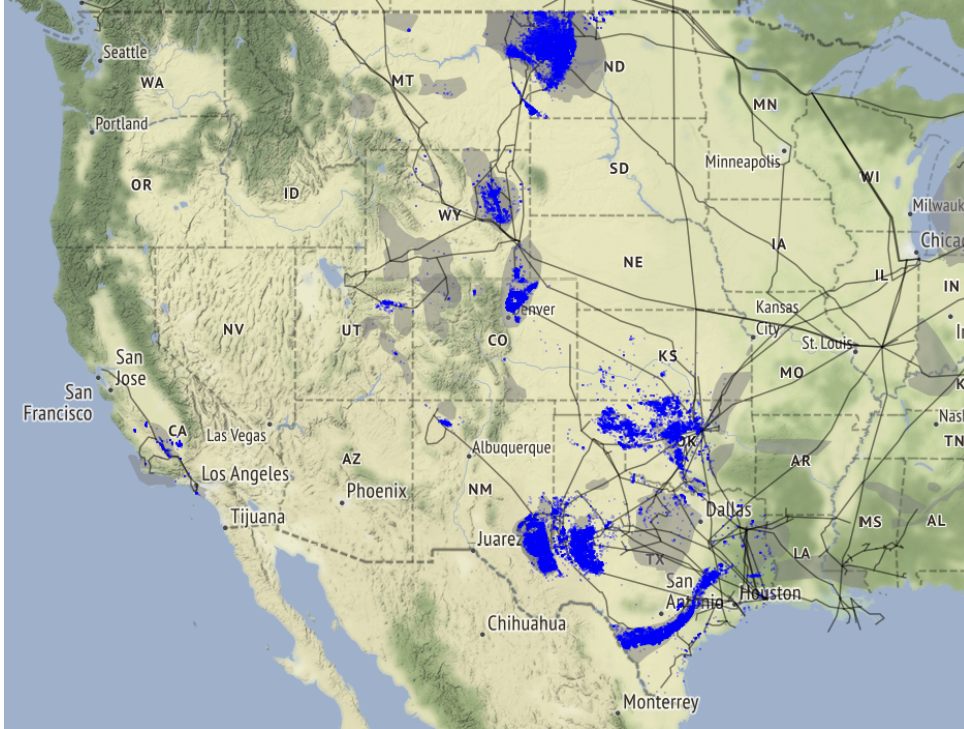


Figure 1. Plot of horizontally drilled wells contained in the data set. Blue dots mark the locations of horizontally drilled wells. Grey patches correspond to shale formations and black solid lines are crude oil pipelines. Shape files for shale formations and pipeline infrastructure are provided by the U.S. Energy Information Administration. Map is constructed using OpenStreetMap.

given month, we observe several well-specific time-invariant characteristics for each well. This includes, but is not limited to, well location, well operator and well drilling direction. Regulatory reporting standards distinguish between vertical, directional and horizontal drilling directions. For our analysis, we identify unconventional wells that require hydraulic fracturing to be those that are horizontally drilled.

Figure 2 shows examples of typical oil well production profiles. Panels A and B are examples of two shale wells. Panels C and D show conventional wells. All four wells are located in Texas. The four panels indicate several differences between horizontally and vertically drilled wells. First, the initial production month does not correspond to peak output. During this period, the operator performs a test-run which do not necessarily reflect the productivity potential of a given well. We should therefore exclude the first production month from our well-level analysis. Second, shale wells are typically more productive than their conventional counterparts with considerably larger peak outputs. Combined with rapid decline rates, this means that unconventional wells have more front-loaded production profiles. Panels A and B also show examples of the aforementioned shale well-specific re-stimulation behaviour. When output becomes relatively low, the operator exercises her option to re-fracture the well to increase well productivity. The need for hydraulic fracturing in the completion stage, front-loaded production profiles and

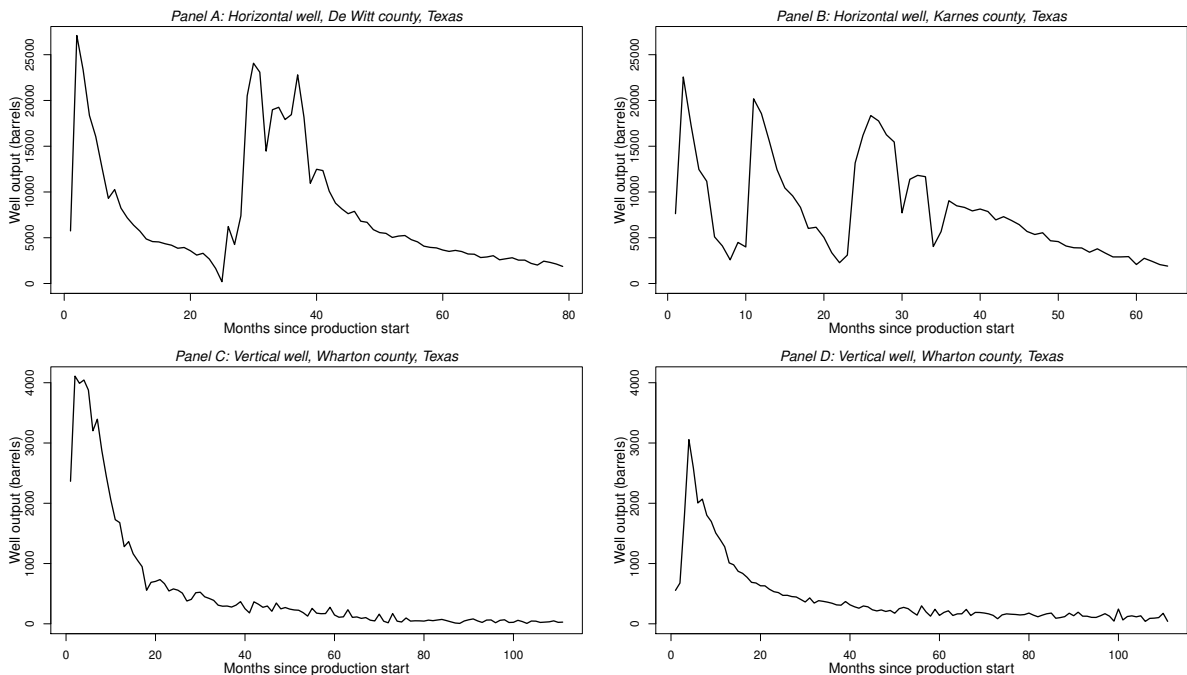


Figure 2. Examples of well production profiles for different production technologies. *Horizontal* refers to shale wells and *vertical* to conventional wells.

the option to re-stimulate wells are all specific traits to unconventional wells that suggest producers can move output inter-temporally to optimize expected profit.

To see how these traits generalize in the cross-section, Figure 3 shows the mean output across conventional and unconventional wells for the same production month in a well life-cycle. As is evident, the central tendency is that shale wells are more productive by several orders of magnitude and thus can yield significant revenue immediately following well completion. We also note that there does not seem to be any indication that there is a relationship between well age and the decision to re-stimulate wells in the cross-section. If that was the case, we should have seen that the average well output increased systematically with well age. If anything, there is a small hump between months 80 and 160, but we deem it to be most likely a by-product of the sample size shrinking as well-age increases. This suggest that well-age in itself is not sufficient to explain why refracturing events take place. To investigate further, we report the full distributions of log well output at four different points in the well life cycle, the 1st, 40th, 80th and 120th production months in Figure 4. While the figure show considerable heterogeneity in the production level of wells at various horizons, there is a pronounced shift to the left in the distribution at the 40th production month compared to the distribution at the 1st production month. In fact, only the very lower right tail of the 40th production month covers the median production level at the 1st production month, indicating that refracturing events are most likely to occur during the first three years a well is active. We also note that the production distributions continue to shift to the left at as the horizon

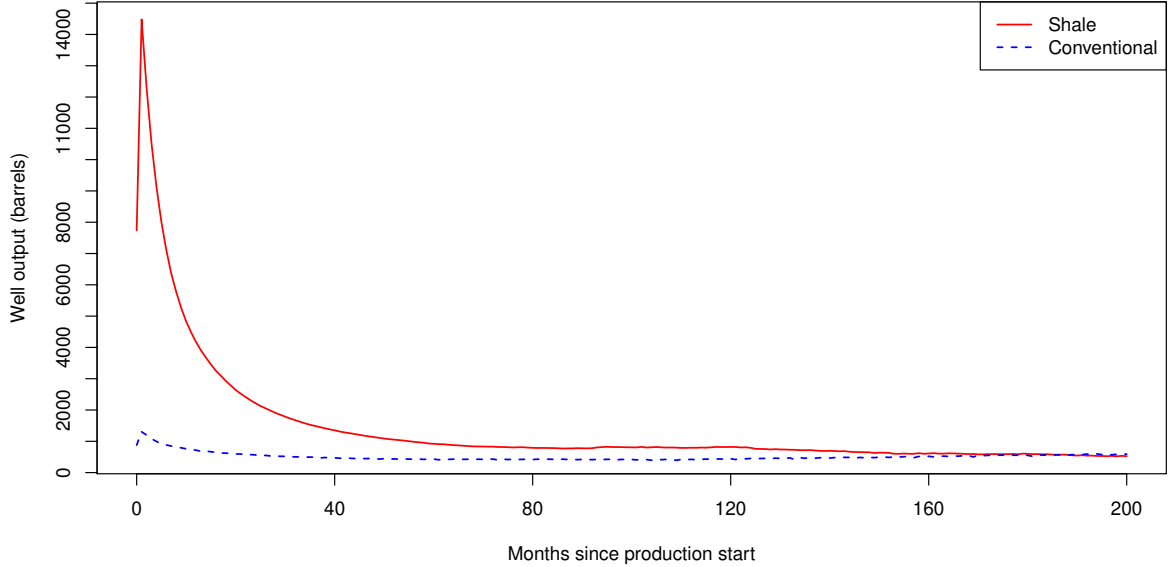


Figure 3. Mean production profiles of conventional and shale wells located in Texas. Each data point is computed by taking the mean output across all wells at the same point in their respective life-cycles. Only wells that began production between 2005:M01 and 2017:M12 are included in the computations.

increases, but the shifts are less pronounced for the larger horizons. Finally, we observe that the data indicate a skewed left tail distribution for all horizons. Especially for the 80th and 120th production months, we observe a long left tail that is almost without mass past 3, indicating some outlier observations.

The ability of shale well operators to stimulate uncompleted wells or re-stimulate producing wells in a timely manner hinges on fracking contractors and supplies being readily available. There could be circumstances, e.g. when expected future market conditions look favourable, when demand for input factors other than water like fracking crews, equipment and supplies outstrip availability. One would then expect that areas where one finds the highest geographical density of shale wells to have the least severe bottlenecks and thus the largest responses to price swings. Examples of such areas are the most developed shale plays in our sample such as the Eagle Ford and Permian plays in Texas, the Bakken play in North Dakota, and the Mississippian and Anadarko plays in Oklahoma. There is another consideration that oil producers have to make when deciding to refracture a well. Because shale wells are drilled horizontally, they can inflict negative externalities when in close proximity to each other. In particular, the fracking process can create cracks unpredictably and if the well-spacing is small, the refracturing of one well might interfere with the production potential of surrounding wells. We will thus expect wells that are located closely together, say less than 600 feet, to be less responsive to oil prices on average. As more wells are developed in a given area, well-spacing tends to

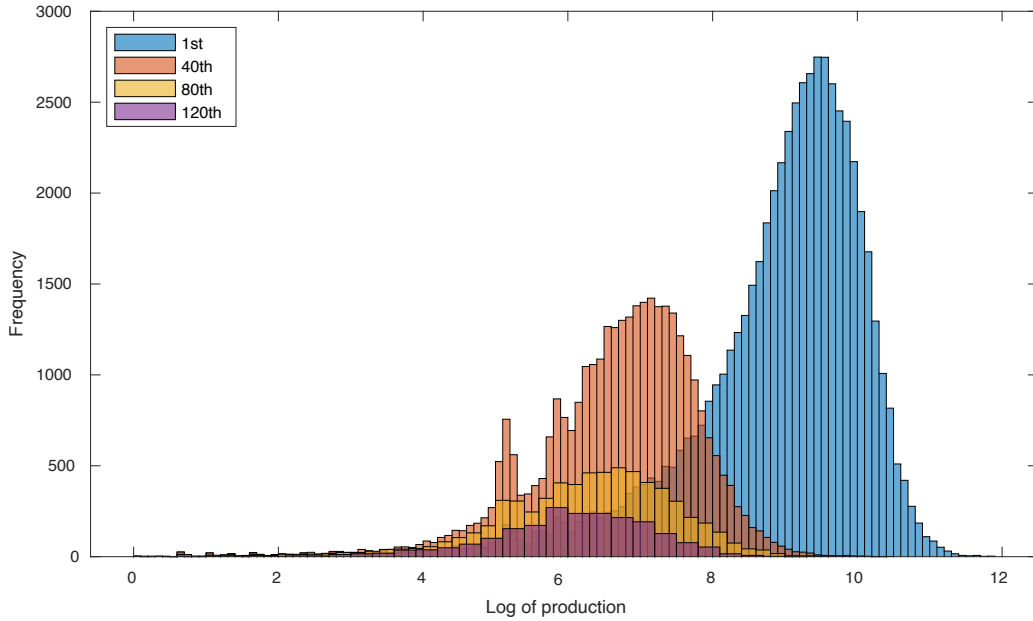


Figure 4. Distribution of the log of well output at the 1st, 40th, 80th and 120th production months across all shale wells in the dataset.

become tighter over time.

We impose some minor sample selection designed to reduce errors in the computation of price responsiveness of oil producers. First, there are wells that have one or more episodes of producing zero barrels during a month either from being shut in, producing only natural gas or through reporting errors.⁵ An operator can choose to shut down a well in preparation for a natural disaster or when they deem production to be too costly relative to the current market price. In our data, a well can have a reported output of zero, but be producing natural gas in those periods. We exclude such wells from our analysis because they create extreme outliers. Hence, we drop all wells that have recorded zero barrels produced in at least one month from the sample. Second, the first recorded production data point per well is unlikely to account for a full month and is considered test production. Output in this phase does not necessarily reflect the ultimate productivity of the well (as can be seen in Figure 2) and is therefore dropped. The resulting unbalanced panel runs over the time period 2005:M01–2017:M12 and encompasses 83,244 shale wells. Of these, 50% are located in the state of Texas and 20% in North Dakota. Texas dominates due to its size, geology and mature oil sector. North Dakota became a major oil producer following the advent of hydraulic fracturing in the Bakken shale play. Colorado, New Mexico and Oklahoma have about 8% each, while California, Kansas, Montana, Utah and Wyoming share the rest. It is important to stress however, that

⁵In at least some jurisdictions, contracted liquids gatherers and not the producers themselves are responsible for reporting well production to the state regulatory agencies from which our data is based.

not every state is made alike. California for instance, is a minor producer of shale oil separated topographically to the other states by the Rocky Mountains. Of the states in our sample, North Dakota, Oklahoma and Texas are the ones understood to be major shale oil producing states. In recent years, the Permian *shale play* located in New Mexico has also been developed. After having applied the same data cleaning procedure describe above, our sample of conventional wells in Texas consists of 87,963 wells.

3 Models and results

Our aim is to study the price-responsiveness of U.S. shale oil producers. We start by estimating the aggregate response of U.S. shale producers to oil prices using aggregate data. We then show how the aggregate price-responsiveness change when using a microdata panel at the well level. Finally we extend the analysis by exploring in the cross-section, e.g. across states and with quantile regressions across the well output distribution.

3.1 Responsiveness of aggregate shale producers

Our starting point is a standard oil supply equation.

$$\ln q_t = \mu + \eta_{oil} \ln P_t^{oil} + X_t + \varepsilon_t \quad (1)$$

where q_t is aggregate oil production, P_t^{oil} is the WTI spot price, X_t is a vector of macro controls and μ is a constant. The aggregate price responsiveness of U.S. oil producers, the short-term supply elasticity, is measured by η_{oil} and can be estimated. However, there are numerous concerns with estimating Equation 1 and thereby η_{oil} . First, it suffers from a reverse causality problem, as aggregate U.S. oil production is likely to impact oil prices, causing a simultaneity bias in the estimated η_{oil} . Therefore, a common assumption in the literature when identifying for instance oil market structural VAR models is that oil producers cannot adjust their production within a month in response to price shocks, see e.g., Kilian (2009).⁶ Such an assumption is supported by results on conventional U.S. oil producers in earlier empirical studies, i.e., Pesaran (1990), Dahl and Yucel (1991), Ramcharran (2002), Smith (2009), and Anderson, Kellogg, and Salant (2018). In particular, Anderson, Kellogg, and Salant (2018) find no evidence of Texan conventional oil wells adjust production as the price of oil changes. Production from existing wells is instead constrained by reservoir pressure, which decays slowly and steadily as oil is extracted. Based on these results, a small supply elasticity is justified due to long lead-times in well

⁶Alternatively, the short-term supply elasticity is assumed to be bounded between zero and a very small number, see Kilian and Murphy (2012, 2014).

development. It is therefore common to include lagged oil prices instead of contemporaneous oil prices as regressors. However, recently this assumption have been called into question by Baumeister and Hamilton (2019), Caldara, Cavallo, and Iacoviello (2019) and Bornstein, Krusell, and Rebelo (2021). Baumeister and Hamilton (2019) use a flexible Bayesian identification approach to incorporate uncertainty regarding the value of the supply elasticity and that allows for an simultaneous response between oil prices and oil production. Caldara, Cavallo, and Iacoviello (2019) calculate the supply elasticities by minimizing the distance between the elasticities found in structural VARs and some selected target values constructed from a survey of relevant studies (i.e., Hamilton (2009)). Bornstein, Krusell, and Rebelo (2021) argue that fracking allows for a more flexible production process than conventional oil production, which enables shale well operators to potentially reduce extraction rates when market conditions are poor, or resume extraction when conditions are improving.

A second challenge with estimating an aggregate oil production equation is that it may suffer from aggregation bias. For instance, aggregating over all individual wells is equivalent to imposing identical parameter values for all producing wells regardless of well, geology, technology or firm characteristics. Such bias may be particularly large when there are heterogeneity in the cross-section. A more suitable approach is therefore to study the price-responsiveness of oil producers using micro data at the well level. For instance, we know that there can be large differences in the initial production level across wells, possibly reflecting unique geological factors.

3.2 Well panel analysis

We use a unique well-level panel data set at the monthly frequency covering all reported onshore shale oil wells producing in the ten major oil-producing states. In addition, we also have data for conventional wells in Texas. While our data set is more comprehensive than earlier micro studies in the literature, we are not the first to analyze the price-responsiveness of U.S. oil producers using microdata. As discussed in the introduction, there are two related studies, Bjørnland, Nordvik, and Rohrer (2021) and Newell and Prest (2019). However, the two studies reach opposite conclusions. While Bjørnland, Nordvik, and Rohrer (2021) find that shale wells respond positively to oil price increases, Newell and Prest (2019) find no evidence that shale producers respond on impact to oil price increases. The two studies, however, differ both in terms of their data sets as well as their econometric model specification. There are five important differences between the two models. First, Bjørnland, Nordvik, and Rohrer (2021) use data for shale wells in North Dakota, while Newell and Prest (2019) use data from shale wells in five U.S. states (including North Dakota). Second, Bjørnland, Nordvik, and Rohrer (2021) estimate the

model in first differences, while Newell and Prest (2019) estimate their model in log-levels but includes a cubic spline to account for the typical production profile of a shale well. Third, Bjørnland, Nordvik, and Rohrer (2021) estimate a model with well age and well fixed effects, while Newell and Prest (2019) instead include only well fixed effects. Fourth, Bjørnland, Nordvik, and Rohrer (2021) include both the spot price and the spot-futures spread in their model, arguing that the latter carries important information about producers' price expectations. Newell and Prest (2019) do not include the spot-futures spread in their model, but instead argue that it is important to control for natural gas prices, since many oil producers are also gas producers. Finally, the two papers also differ in terms of their controls. Bjørnland, Nordvik, and Rohrer (2021) include, well and year fixed effects in addition to adding several macroeconomic controls, while Newell and Prest (2019) only include well fixed effects.⁷

Being aware of the two different model specifications, we construct a baseline model which covers the most important features of both models. We estimate the following baseline model

$$\begin{aligned} \ln q_{it} = & \eta_{oil} \ln P_t^{oil} + \eta_F (\ln P_t^{oil} - \ln F_{t,t+3}) + \eta_{gas} \ln P_t^{gas} \\ & + X_t + g(\text{Age}_{it}) + \lambda_y + \mu_i + \varepsilon_{it} \end{aligned} \quad (2)$$

where q_{it} is oil production in terms of barrels for well i at time t . On the right hand side, P_t^{oil} is the WTI spot price, $F_{t,t+3}$ is the 3-month WTI futures price and P_t^{gas} is the Henry Hub natural gas spot price. We construct these time series as a monthly averages over daily observations. $g(\text{Age}_{it})$ is a cubic spline with knots at every twelfth month constructed from the age of the individual well, λ_y is a year fixed effect and X_t is a vector of macro controls. This vector consist of the federal funds rate, the copper price, the U.S. foreign exchange rate, the MSCI world stock index and the VIX index. These are included to control for aggregate demand and uncertainty in financial markets. We apply two-way clustering across well-time when computing the standard errors to allow for within-well time-dependence as well as within year-month across well correlation in the error term. The latter will allow for correlation across wells that occur when wells respond to the same oil price movements. Since the main source of variation in this study, the oil price, is common to all individual wells and there is strong serial dependence for each well, this feature is critical for valid inference. We are interested in the sum $(\eta_{oil} + \eta_F)$ and the appropriate standard error. To obtain these estimates, we run an auxiliary regression

⁷Newell and Prest (2019) estimate two separate specifications. The first includes technical details about the well such as well depth as control variables. The second includes only well fixed effects. It is the latter that we consider for our comparison.

where we have added and subtracted $\eta_F(\ln P_t^{oil})$ from Equation 2.⁸

We start with reporting results from models estimated on aggregate data for unconventional and conventional oil production, respectively. To facilitate direct comparison with results using well-level information, we estimate an aggregate version of Equation 2, including the same regressors and controls. Results for specifications that are estimated in log-levels are shown in columns (1) and (3) for shale and conventional production, respectively, while similar results for specifications in log-differences are reported in columns (5) and (7). We find no evidence of an effect of the spot price and spot-futures spread on both aggregate shale and aggregate conventional oil production.⁹

As discussed above, a challenge with estimating an aggregate oil production equation is that it may suffer from both simultaneity bias and aggregation bias. We therefore now turn to study the price-responsiveness of oil producers using a well-level panel, as described in Equation 2. In Table 1, columns (2) and (6), we report results for estimating a model with shale wells in log-levels and in log-differences, respectively. Consistent with Bjørnland, Nordvik, and Rohrer (2021) and Newell and Prest (2019), we find a negligible effect of the spot price on aggregate shale oil production. However, we complement Bjørnland, Nordvik, and Rohrer (2021) by finding a significant and strong response of the spot-futures spread, which is estimated to be 0.68, yielding a total price-responsiveness of around 0.6. Results remain robust when specifying the model in first differences (column 6), although the coefficient increases somewhat. This shows that the spot-futures spread carries important forward looking information which U.S. shale producers reacts to. When the spot-futures spread increases, the oil spot price is expected to fall in the future, urging shale producers to increase their production now. Importantly, we see that we can reconcile the opposite conclusions reached by Bjørnland, Nordvik, and Rohrer (2021) and Newell and Prest (2019) by introducing the spot-futures spread. As expected, we do not find a similar price responsiveness for conventional oil wells. Results in Table 1, columns (4) and (8), show negligible effects of both the spot price and spot-futures price on aggregate conventional production. This reaffirms results in Anderson, Kellogg, and Salant (2018) and mirrors the results of Bjørnland, Nordvik, and Rohrer (2021) (for North

⁸Performing this transformation can be problematic as the correlation between the spot and future prices is large and can cause multicollinearity. While this may be an issue for the estimation of the auxiliary model, it is not a cause for concern in the original specification as the spot price and the spot-futures spread are not highly correlated. Thus the separate estimates of η_{oil} and η_F in the non-transformed model can be reliably obtained. The fact that the separate estimates of η_{oil} and η_F versus $(\eta_{oil} + \eta_F)$ are numerically identical suggests that multicollinearity does not bias the estimates of the auxiliary model. If anything, we should worry about the standard errors being too large.

⁹While not reported here, results are also similar if we instead used the lagged value of spot prices and spot-futures spread, commonly used in the literature.

Shale vs. conventional wells on aggregated and panel data

Specification	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	$\ln q_t$				$\Delta \ln q_t$			
	Shale		Conventional		Shale		Conventional	
η_{oil}	0.02 (0.06)	-0.06* (0.04)	-0.02 (0.04)	-0.01 (0.03)	0.09 (0.06)	-0.15 (0.09)	-0.01 (0.05)	-0.05 (0.06)
η_F	-0.52 (0.41)	0.68*** (0.25)	-0.10 (0.25)	-0.11 (0.22)	-0.08 (0.23)	0.96** (0.42)	0.12 (0.26)	0.36 (0.30)
η_{gas}	-0.03 (0.04)	-0.03 (0.03)	0.00 (0.02)	-0.01 (0.02)	-0.03 (0.03)	-0.03 (0.04)	-0.01 (0.03)	0.00 (0.03)
$\eta_{oil} + \eta_F$	-0.50 (0.19)	0.62*** (0.23)	-0.12 (0.23)	-0.11 (0.21)	0.01 (0.19)	0.82** (0.36)	0.11 (0.24)	0.31 (0.27)
Macro controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Well FE	No	Yes	No	Yes	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Linear trend	Yes	No	Yes	No	Yes	No	Yes	No
Well Age FE	No	Spline ³	No	Spline ³	No	Spline ³	No	Spline ³
First observation	2005:M01	2005:M02	2005:M01	2005:M02	2005:M01	2005:M02	2005:M01	2005:M02
Last observation	2017:M12	2017:M12	2017:M12	2017:M12	2017:M12	2017:M12	2017:M12	2017:M12
N		58,422		84,760		57,455		80,653
$N \times T$	156	2,649,951	156	5,700,878	156	2,543,410	156	5,238,997
\bar{R}^2	0.72	0.77	0.85	0.82	0.44	0.08	0.40	0.13
Clustering	Time	Well-Time	Time	Well-Time	Time	Well-Time	Time	Well-Time
Num. clusters	156	155	156	154	156	154	156	153

Table 1. Estimation results for various baseline model specifications. η_{oil} is the coefficient on the WTI spot price, η_F is on the spot-futures spread, η_{gas} is on the Henry Hub gas spot price and the coefficient ($\eta_{oil} + \eta_F$) is the sum of the spot and spread coefficients estimated by an auxiliary regression. Columns 1–4 are for models in log-level and 5–8 on log-difference. Columns 1, 2, 5 and 6 are for shale wells while columns 3, 4, 7 and 8 are on vertically drilled Texas wells. Columns 1, 3, 5 and 7 are on aggregated data and columns 2, 4, 6 and 8 are on our well-level panel. The macro controls consist of the federal funds rate, the trade-weighted foreign exchange rate, the copper price, the MSCI world stock index and the VIX index. The cubic spline has knots at every 12th production month. N refers to the number of unique wells included in the estimation.

Dakota) and Newell and Prest (2019).

3.3 The price responsiveness of U.S. shale oil producers

Having established a significant and positive price responsiveness of U.S. shale producers, we are interested in further uncovering what characterizes their responsive behaviour. As a starting point, we investigate whether the price responsiveness depends on the production level. There are at least two reasons why the production level could affect the price responsiveness of a well. First, as discussed in Section 2, shale wells have front-loaded production profiles. One would therefore expect there to be a positive association between

favourable oil prices and output given that the production level is high during this phase of a well life-cycle. This phase is initiated by a *well completion* event and as such, a possible interpretation can be that, on average, well completions are associated with favourable market conditions. This would favour a large production response when oil prices are high and expected to fall. For wells that are at a later stage in their life-cycles with lower output levels, we would generally expect a more muted price response. Finally, if the oil producer believes there is an additional revenue potential in a well, the producer may be more willing to take on the additional cost of *refracturing* the well. Once a well is refractured, the production can increase substantially—in many cases almost to the same level as initial production levels. This would indicate a strong price response when the production level is high.

We study whether the price responsiveness depends on the production level, by estimating a quantile regression version of Equation 2.¹⁰ Since we are interested in the marginal effects of the spot price and spot-futures price on shale production, we use the unconditional quantile regression approach developed by Firpo, Fortin, and Lemieux (2009).¹¹ This method consists of running a regression of the recentered influence function (RIF) of the unconditional quantile on the explanatory variables. This allows for a marginal effect interpretation similar to the one of standard OLS. In implementing the Firpo, Fortin, and Lemieux (2009) approach in our panel data setting with high-dimensional fixed-effects and multi-way clustering of standard errors, we use the Stata code developed by Rios-Avila (2019). We focus on the 1st, 5th, 25th, 50th, 75th, 95th and 99th percentiles.¹²

Table 2 shows the quantile estimation results. The line *Mean LHS* gives an indication of where in the output distribution the different percentiles are found. The median level of well production is $\exp(7.26) = 1422$ barrels per month. This may seem small for a shale well, but Figure 3 indicates that the average shale well reaches this production level within 40 months after their initial start of production month. At the median, there is

¹⁰Quantile regressions have recently been a popular tool for studying drivers of macroeconomic tail risks. For instance, Adrian, Boyarchenko, and Giannone (2019) study the conditional distribution of GDP growth as a function of economic and financial conditions and argues that financial conditions are particularly informative about future downside macroeconomic risk.

¹¹It is more common to estimate conditional quantile regressions. However, this would mean estimating the marginal effect conditional on the values of all other covariates that are included in the model. Given the considerable amount of control variables that we include in our regressions, it would be unappealing to pick specific values of these variables when computing the price responsiveness of shale producers.

¹²To control for substantial one-off output dips (outliers) in otherwise ordinary well production profiles, we include dummy variables that are equal to 1 if such a dip event occurs for well i at time t . As is indicated by Figure 4, the output distribution has a long left tail and the results in the lower percentiles may be sensitive to these outlying observations.

Unconditional quantile regression on log-level full cross-section

Distributional stat.	Q1	Q5	Q25	Q50	Q75	Q95	Q99
η_{oil}	0.16 (0.10)	0.08** (0.04)	-0.01 (0.04)	-0.02 (0.04)	-0.10** (0.04)	-0.21** (0.10)	-0.25* (0.15)
η_F	0.93 (0.62)	0.10 (0.18)	0.37 (0.24)	0.10 (0.27)	0.48* (0.28)	1.97*** (0.47)	2.31*** (0.63)
η_{gas}	-0.01 (0.05)	-0.03* (0.02)	-0.01 (0.02)	-0.02 (0.03)	-0.02 (0.03)	-0.05 (0.04)	0.00 (0.06)
$\eta_{oil} + \eta_F$	1.09 (0.68)	0.19 (0.16)	0.35 (0.22)	0.09 (0.25)	0.38 (0.26)	1.76*** (0.40)	2.06*** (0.52)
Macro controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Well FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Well Age FE	Spline ³	Spline ³	Spline ³	Spline ³	Spline ³	Spline ³	Spline ³
Dip dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
First observation	2005:M02	2005:M02	2005:M02	2005:M02	2005:M02	2005:M02	2005:M02
Last observation	2017:M12	2017:M12	2017:M12	2017:M12	2017:M12	2017:M12	2017:M12
N	58,422	58,422	58,422	58,422	58,422	58,422	58,422
$N \times T$	2,649,951	2,649,951	2,649,951	2,649,951	2,649,951	2,649,951	2,649,951
\bar{R}^2	0.62	0.40	0.58	0.62	0.57	0.38	0.26
Mean LHS	3.10	4.94	6.37	7.26	8.07	9.26	9.96
Clustering	Well-Time	Well-Time	Well-Time	Well-Time	Well-Time	Well-Time	Well-Time
Num. clusters	155	155	155	155	155	155	155

Table 2. Unconditional quantile regression estimation results on data in log-levels. Estimation is based on Firpo, Fortin, and Lemieux (2009). Parameters η_{oil} and η_{gas} are the coefficients on the natural log of WTI and Henry Hub spot prices. η_F is the coefficient on natural log of the spot-futures spread. $(\eta_{oil} + \eta_F)$ is the total response of quantities produced from the level of spot-futures spread estimated by an auxiliary regression. *Mean LHS* gives an indication on where in the distribution each percentile is located. All wells are shale wells. The dip dummies are included to control for sudden one-off dips in otherwise ordinary production profiles that cause outliers. For example, we construct a dummy that is equal to 1 for well i at time t if output at time t is below 25 barrels and the observations for $t - 1$ and $t + 1$ are larger than 25. 25 barrels per month is an extraordinarily small amount of output for a shale well.

no statistically significant association with the oil prices. However, that is not the case in the upper tail of the distribution. With all our controls and fixed effects included, we find that there is a strong association between production output and oil prices given that the production level is high. The upper section of the production distribution is likely to pick up either wells that have just started their production (been completed) or wells that have just been refractured. Given the front-loaded production profile of shale wells, it is reasonable that producers would like to take advantage and there being an association

between high prices and high production levels.

Regression results conditional on refracturing and production start				
Specification	(1)	(2)	(3)	(4)
	$\ln(q_{it})$	$\ln(q_{it})$	$\ln(q_{it})$	$\ln(q_{it})$
η_{oil}	-0.06*	-0.06*	-0.06*	-0.06*
	(0.04)	(0.4)	(0.04)	(0.04)
$\eta_{oil} \times (start_{it} = 1)$		0.07***		0.07***
		(0.01)		(0.01)
$\eta_{oil} \times (refractured_{it} = 1)$			-0.19***	-0.20***
			(0.01)	(0.01)
η_F	0.68***	0.66***	0.67***	0.65***
	(0.25)	(0.24)	(0.24)	(0.24)
$\eta_F \times (start_{it} = 1)$		0.26		0.25
		(0.39)		(0.39)
$\eta_F \times (refractured_{it} = 1)$			1.90**	1.93**
			(0.76)	(0.75)
η_{gas}	-0.03	-0.03	-0.03	-0.03
	(0.03)	(0.03)	(0.03)	(0.03)
Macro controls	Yes	Yes	Yes	Yes
Well FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Well Age FE	Spline ³	Spline ³	Spline ³	Spline ³
First observation	2005:M02	2005:M02	2005:M02	2005:M02
Last observation	2017:M12	2017:M12	2017:M12	2017:M12
N	58,422	58,422	58,422	58,422
$N \times T$	2,649,951	2,649,951	2,649,951	2,649,951
\bar{R}^2	0.77	0.77	0.78	0.77
Clustering	Well-Time	Well-Time	Well-Time	Well-Time
Num. clusters	155	155	155	155

Table 3. Estimation results on the full cross-section of our dataset with data on log-levels. Parameters η_{oil} and η_{gas} are the coefficients on the natural log of WTI and Henry Hub spot prices. η_F is the coefficient on natural log of the spot-futures spread. $refractured_{it}$ is a dummy variable equal to 1 if the well is likely to have been refractured at time t . $start_{it}$ is a dummy variable equal to 1 if t is the first full production month for well i . All wells are shale wells.

To further explore this, we add interaction effects to our baseline model in Equation 2. We construct two dummy variables to account for the front-loaded production. $start_{it}$ is a dummy variable which equals 1 if t is the first full production month for well i and $refractured_{it}$ is a dummy variable which equals 1 if the well is likely to have been refractured at time t . The variable $refractured_{it}$ is set to equal 1 if the level or production between period t and $t + 1$ increases by 2000 barrels or more.¹³ We report results in Table 3. In column 1 we list the results for our baseline model without interaction terms.

¹³Our results are robust to setting the number of barrels to for instance 1000 or 5000. Results for the

In column 2 we show results for a model where we add the interaction dummy $start_{it}$. Somewhat surprisingly, the additional impact of the interaction terms is small, indicating no additional effect for newly started wells. The results are however more supportive when adding the interaction dummy $refractured_{it}$. As shown in column 3 in Table 3, both interaction terms are statistically significant, indicating that the price response indeed is stronger for wells going through a refracturing event.

Together with the results reported from the unconditional quantile regression analysis, we conclude that there is substantive evidence in the data that unconventional oil wells respond systematically to signals of favourable oil prices. Furthermore, these responses are found to be associated with the two ways unconventional oil producers can time their production decisions: well completion and refracturing.

3.4 Well ownership, well spacing and a firm-level panel

We continue our analysis by exploring three additional sources of heterogeneity. First, we study whether the price-responsiveness varies across the distribution of well ownership. We hypothesize that larger firms have more resources available to quickly act on signals of favourable market conditions. Second, we want to learn whether there is evidence of a well spacing externality. That is, if refracturing of a well located in close proximity to other wells may cause adverse effects on the production potentials of those wells. And third, we recognize that the relevant decision-maker is the individual firm. We therefore study whether firms expand their production volumes when prices are expected to increase.

To investigate whether well ownership affects the price-responsiveness, we proceed by counting the number of wells operated by each firm. We then examine the distribution of well counts and identify the firms that together account for 25% of the wells in operation in our sample. These firms are in descending order EOG Resources, Occidental Petroleum (Oxy), Chesapeake, Marathon Oil, Continental Resources, ExxonMobil and ConocoPhillips. Using this information, we construct a dummy variable $large_i$ that is equal to 1 if well i is operated by one of these firms. To assess the additional price-response for wells operated by these seven firms relative to the other firms, we add interactions between the price variables and the dummy variable. The estimation results are reported in Panel A of Table 4.

The results show that there is an additional statistically significant average response to the spot-futures spread for wells that are operated by the seven largest firms. In fact, the estimates suggest that the total price responsiveness effect is more than twice as large as for the other firms. The total effect ($\eta_{oil} + \eta_F$), including the interaction effect, is 0.95 and interaction effect with the dummy $start_{it}$ is also robust to including more than just the starting month. We have checked results when including up to the first six months of production.

Panel A: Firm size and price response		Panel B: Well spacing and price response		Panel C: Price response in firm-level panel	
	$\ln(q_{it})$		$\ln(q_{it})$		$\ln(\tilde{q}_{kt})$
η_{oil}	-0.03 (0.04)	η_{oil}	0.11** (0.04)	η_{oil}	-0.20*** (0.06)
$\eta_{oil} \times (large_i = 1)$	-0.11*** (0.02)	$\eta_{oil} \times (space_i = 1)$	-0.25 (0.03)		
η_F	0.52** (0.25)	η_F	-1.04*** (0.39)	η_F	0.86** (0.34)
$\eta_F \times (large_i = 1)$	0.57** (0.23)	$\eta_F \times (space_i = 1)$	2.44*** (0.39)		
η_{gas}	-0.03 (0.03)	η_{gas}	-0.03 (0.03)	η_{gas}	-0.01 (0.04)
$(\eta_{oil} + \eta_F)$	0.49** (0.23)	$(\eta_{oil} + \eta_F)$	-0.93** (0.36)	$(\eta_{oil} + \eta_F)$	0.65*** (0.25)
$(\eta_{oil} + \eta_F) + (\eta'_{oil} + \eta'_F) \times (large_i = 1)$	0.95*** (0.28)	$(\eta_{oil} + \eta_F) + (\eta'_{oil} + \eta'_F) \times (space_i = 1)$	1.26*** (0.25)		
Macro controls	Yes	Macro controls	Yes	Macro controls	Yes
Well FE	Yes	Well FE	Yes	Firm FE	Yes
Year FE	Yes	Year FE	Yes	Year FE	Yes
State FE	Yes	State FE	Yes	State FE	Yes
Well Age FE	Spline ³	Well Age FE	Spline ³	State-Firm trend	Linear
First observation	2005:M01	First observation	2005:M01	First observation	2005:M01
Last observation	2017:M12	Last observation	2017:M12	Last observation	2017:M12
N	58,422	N	57,273	N	1,050
$N \times T$	2,649,951	$N \times T$	2,597,681	$N \times T$	98,910
\bar{R}^2	0.77	\bar{R}^2	0.77	\bar{R}^2	0.69
Clustering	Well-Time	Clustering	Well-Time	Clustering	Firm-Time
Num. clusters	155	Num. clusters	155	Num. clusters	156

Table 4. Panel A shows estimation results for a well-level model where we have interacted the oil prices with a dummy variable for whether the well is owned by a large oil firm. In particular, we use the distribution of well ownership to identify the firms in the dataset that are among the top 25% in term of number of wells in operation. These are EOG Resources, Occidental Petroleum (Oxy), Chesapeake, Marathon Oil, Continental Resources, ExxonMobil and ConocoPhillips. If well i is operated by on of these firms, the dummy variable $large_i$ is equal to 1. Parameters η_{oil} and η_{gas} are the coefficients on the natural log of WTI and Henry Hub spot prices. η_F is the coefficient on natural log of the spot-futures spread. $(\eta_{oil} + \eta_F)$ is the total price response estimated by an auxiliary model. $(\eta'_{oil} + \eta'_F)$ is the total additional price response from the interaction terms. Panel B shows estimation results from a well-level model where we have interacted the oil prices with dummy variable $space_i$ which is equal to 1 if well i is located more than 600 feet (183 metres) away from the closest completed well. Wells that we do not have well spacing information for have been omitted from the estimation sample. Panel C shows estimation results from a firm-level model. The model has the same specification as the baseline model, but \tilde{q}_{kt} is the aggregate production volumes across all i wells operated by firm k .

significant at the 1% level. A comparison of this result with those of column 2 in Table 1 indicates that it is the largest firms that respond the most since the baseline average response without the interaction term is 0.62. This aligns with the findings of Bjørnland, Nordvik, and Rohrer (2021) who identifies a stronger effect for the top 99 firms in North Dakota in terms of production volumes.

We now turn to whether spacing between wells affect their price-responsiveness. As

previously discussed, a refracturing event can occur when an operator has favourable expectations about future market conditions and is willing to incur the cost of repeating the process of pumping water, sand and other chemicals into the well to create new or expand existing cracks in the shale rock layer. This event will potentially increase the output of a well by many orders of magnitude. However, the process of refracturing is unpredictable and if wells are not sufficiently spaced apart, proceeding with the refracturing process for a given well can cause irreversible damage to the well and other wells located around it. As such, we should expect oil producers that have drilled wells too close together to be more restrictive at re-stimulating their wells—or put differently, wells located further apart from other wells should respond more to favourable movements in oil prices. From Rystad Energy, we have obtained categorical data on the distance from a given well to the closest neighbouring completed well.¹⁴ The distances are measured from the bottom of the wells in feet. We construct a dummy variable $space_i$ that takes the value 1 if well i is not located within 600 feet (approximately 183 metres) of another well. We interact this dummy variable with the spot price and the spot-futures spread to estimate the additional effect of a well being spaced more than 600 feet apart from its closest neighbour.

The estimation results are reported in Panel B of Table 4. Also in this case, there is a significant positive additional effect from the interaction term on the spot-futures spread. This means that wells that are spaced more than 600 feet from their neighbouring wells tend on average to have a stronger response relative to wells that are located closer than this. Approximately 40% of the wells in the sample belong to the ≤ 600 feet category. The total estimate of 1.26 for wells that are spaced more than 600 feet from their neighbouring wells is more than twice as high as the baseline average response without the interaction of 0.62.¹⁵ Furthermore, the estimated response for $space_i = 0$ is negative. We are cautiously interpreting this result as an indication of wells located closely together lose well productivity in response to positive oil price expectations. Put differently, one can make the argument that a refracturing event on average causes irreparable damage to the wells which results in large productivity losses as the results in Table 4 suggest. However, we cannot tell with certainty whether this happens in practice or if the reported results capture some other mechanism related to tight well spacing and production decisions. As the shale patch becomes more developed, well density tends to increase and it becomes more difficult to maintain appropriate well spacing. As a result, one might expect that refracturing events become more infrequent or less effective over time in a given area.

We now turn to study whether firms expand their production volumes when prices

¹⁴Instead of the exact distance from well i to well j , the data places wells into categories by distance. The shortest distance category is defined as wells located ≤ 600 feet apart.

¹⁵Note that the estimation sample when including well spacing information excludes about 1,000 wells that we do not have well spacing data for.

are expected to increase. We proceed by aggregating our well-level data for each firm to obtain a firm-level panel. Since we observe the same firms across multiple states, we include a state fixed effect as well as firm fixed effects to control for unobserved firm- and state-level variation. The standard errors are now clustered on firm-time for the same reasons as previously argued for well-time. We substitute the well-specific cubic spline by a linear trend as we do not expect to see the typical well decline curves for firm-level output. We allow the linear trends to be state-firm-specific. Apart from these modifications, we retain the baseline model specification. The results are reported in Panel C of Table 4. We find the same positive and statistically significant response to oil prices as we found for individual wells. We interpret this as evidence that individual shale well operating producers on average increase their aggregate production across the states in our sample when they receive signals of increasing prices. This is in line with the findings of Bjørnland, Nordvik, and Rohrer (2021).

3.5 Individual states and shale plays

We complete the analysis by examining the price-responsiveness across geographic regions. In particular, we are interested in assessing to what extent the price responsiveness differ across state borders and within the so-called *shale plays*. We go about this by estimating our baseline model separately on subsets of the data by geography.

Table 5 provides information on a variety individual state characteristics. We note that Texas and North Dakota are by far the largest shale oil producing states, followed by Montana and Oklahoma. However, by 2017 all states, with the exception of California, Kansas and Utah, have more than 50% of their total oil production stemming from shale oil producers. Notably, California and Kansas have had a reduction in the share of shale wells between 2014 and 2017. At the same time, Utah saw its share increased, but had no significant reduction in the market share of top five firms. The Utah Herfindahl-Hirschmann Index also remains the highest among these states, which suggests a highly non-competitive environment.

The results for all the states are reported in Table 6. The first column shows the parameter estimates using well level information of the ten oil producing states in our sample taken together, which are identical to those presented in Column 2 in Table 1 above, and included here for comparison. The results from the estimation of the individual states follow in columns 2–11 respectively for California, Colorado, Kansas, Montana, North Dakota, New Mexico, Oklahoma, Texas, Utah and Wyoming.

We find that for seven of the states; Colorado, Kansas, Montana, North Dakota, Oklahoma, Texas, and Utah, there is a significant positive prices response. Of these, all states but Texas and Montana have a response rate that exceeds the average across all ten

Summary of state characteristics

	CA	CO	KS	MT	NM	ND	OK	TX	UT	WY
Share shale 2014	4.19%	84.04%	10.16%	71.62%	59.68%	97.40%	n.a	62.96%	16.05%	44.67%
Share shale 2017	3.68%	93.00%	4.63%	63.31%	78.49%	98.12%	n.a.	75.87%	32.62%	53.64%
HH index full sample	0.31	0.34	0.04	0.14	0.13	0.15	0.14	0.04	0.52	0.22
HH index 2010–	0.31	0.23	0.06	0.13	0.11	0.07	0.06	0.04	0.39	0.16
Market share top 5 firms	87.20%	59.67%	21.04%	71.17%	53.83%	39.45%	21.99%	26.34%	90.27%	45.34%
Market share top 5 firms 2010–	89.52%	83.76%	34.18%	69.41%	61.82%	46.22%	35.12%	36.68%	90.02%	67.22%
Average share of total shale production	0.73%	2.53%	1.59%	12.38%	4.39%	33.10%	4.63%	38.81%	0.24%	1.67%

Table 5. Summary of state characteristics across a variety of dimensions. *Share shale* refers to the share of oil produced in the state that is from shale wells. *HH index* is the Herfindahl-Hirschmann Index which measures market concentration. The higher number for the HH-index, the closer a market is to a monopoly (i.e, the higher the market’s concentration, and the lower its competition). We rank firm size by the number of barrels of crude oil produced. *Total shale production* refers to the total amount of barrels of crude oil produced by shale wells across the ten states. State characteristics data is courtesy of Rystad Energy.

Regression results on log-level state-level data

Subsample	All states	CA	CO	KS	MT	ND	NM	OK	TX	UT	WY
	$\ln(q_{it})$	$\ln(q_{it})$	$\ln(q_{it})$	$\ln(q_{it})$	$\ln(q_{it})$	$\ln(q_{it})$	$\ln(q_{it})$	$\ln(q_{it})$	$\ln(q_{it})$	$\ln(q_{it})$	$\ln(q_{it})$
η_{oil}	-0.06*	-0.04	-0.07	-0.38***	-0.04	-0.11***	0.02	-0.12**	-0.04	-0.06	0.03
	(0.04)	(0.12)	(0.07)	(0.10)	(0.03)	(0.04)	(0.05)	(0.05)	(0.04)	(0.10)	(0.06)
η_F	0.68***	0.17	0.95**	1.47***	0.41**	0.90***	0.52	1.04***	0.51*	1.13*	-0.12
	(0.25)	(0.41)	(0.40)	(0.52)	(0.21)	(0.25)	(0.37)	(0.37)	(0.29)	(0.59)	(0.27)
η_{gas}	-0.03	-0.03	-0.01	-0.09	0.00	-0.07**	0.00	0.02	-0.02	0.06	-0.04
	(0.03)	(0.05)	(0.04)	(0.06)	(0.02)	(0.03)	(0.03)	(0.03)	(0.03)	(0.06)	(0.03)
$\eta_{oil} + \eta_F$	0.62***	0.12	0.88**	1.09**	0.37*	0.79***	0.55	0.92***	0.47*	1.07**	-0.09
	(0.23)	(0.37)	(0.34)	(0.44)	(0.19)	(0.23)	(0.33)	(0.34)	(0.27)	(0.54)	(0.24)
Macro controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Well FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Well Age FE	Spline ³	Spline ³	Spline ³	Spline ³	Spline ³	Spline ³	Spline ³	Spline ³	Spline ³	Spline ³	Spline ³
First observation	2005:M01	2005:M01	2007:M04	2006:M05	2005:M01	2005:M01	2005:M03	2006:M03	2005:M01	2005:M04	2005:M04
Last observation	2017:M12	2017:M12	2017:M12	2017:M12	2017:M12	2017:M12	2017:M12	2017:M12	2017:M12	2017:M12	2017:M12
N	58,422	761	4,286	388	1,255	12,893	3,615	5,526	28,063	219	1,416
$N \times T$	2,649,951	38,759	150,211	13,511	110,560	705,220	161,652	236,259	1,164,466	9,612	59,701
\bar{R}^2	0.77	0.60	0.77	0.77	0.75	0.65	0.78	0.84	0.78	0.82	0.81
Clustering	Well-Time	Well-Time	Well-Time	Well-Time	Well-Time	Well-Time	Well-Time	Well-Time	Well-Time	Well-Time	Well-Time
Num. clusters	155	155	129	139	155	155	153	142	155	153	152

Table 6. Estimation results for each individual U.S. state with data on log-level. Parameters η_{oil} and η_{gas} are the coefficients on the natural log of WTI and Henry Hub spot prices. η_F is the coefficient on natural log of the spot-futures spread. $(\eta_{oil} + \eta_F)$ is estimated by an auxiliary model. All wells are shale wells.

states (0.62). However, for California, New Mexico and Wyoming, the response is smaller and insignificant. California differ from the other ten states by its isolation west of the Rocky Mountains. This, coupled with the minuscule share of shale wells in this state, may explain the low price response. For New Mexico and Wyoming, the explanation is less clear-cut.

New Mexico has a sizeable coefficient estimate that is not statistically significant. The majority of shale production in New Mexico stems from the Permian shale play (see Figure 1). According to our data from Rystad Energy, there was little development of new shale wells in this region prior to the 2014 fall in oil prices, with production growth accelerating as late as 2016. Another possible reason is lack of pipeline infrastructure to carry the crude oil to market in Cushing, Oklahoma.¹⁶ The higher productivity of shale wells can make legacy transportation infrastructure insufficient.¹⁷ These bottlenecks might affect incentives to complete or refracture wells and as such price-responsiveness. In particular, facing local gluts and limited storage capacity, price discounts relative to the WTI spot benchmark will encourage operators to keep the crude oil in the ground. This may explain the lack of response in New Mexico.

Wyoming has, while not statistically significant, a negative price response. While Figure 1 indicates that Wyoming does have legacy infrastructure in place, the same hypothesis of insufficient infrastructure can be posed for this state. Only in the recent years have investments into expanding transportation capacity been prioritized. There is a degree of heterogeneity across states, but apart from the three aforementioned states, we find a positive and in some cases very strong response across the states in our sample.

In our data sample, we have information about whether a well’s geographical location is within a so-called *shale play*—a limited geological shale rock formation where the density of shale wells tends to be higher. A shale play is often located within a larger geological basin with other areas rich in natural resources. While conventional exploration and drilling have been commonplace in these plays for decades, it is mostly following the shale oil boom that these plays have become active due to the low permeability of the reservoirs. Because we possess information about the geographical location of each well, we can use this data to determine whether a particular well is located in a shale play. One should expect the price-responsiveness to be more pronounced in these areas as the local availability of resources and knowledge needed to complete or re-stimulate wells should be higher. We re-estimate our models on subsamples for five major shale plays across three states: the Anadarko and Mississippian plays in the state of Oklahoma; the Bakken play in the state of North Dakota; and the Eagle Ford and Permian plays in the state of Texas. Of these, the latter three shale plays are the most well-known. Bakken and Eagle Ford

¹⁶In its Today in Energy article of November 15th 2017, the U.S. Energy Information Administration writes that “As U.S. crude oil production has increased, particularly in regions such as the Permian basin, so has the need for more transportation infrastructure to accommodate it. However, the rate of production growth and the scale and timing of when additional pipeline capacity is brought online are not always aligned.”

¹⁷Lack of pipeline infrastructure in North America following the shale oil boom and its effects on local prices have been discussed previously by Kilian (2016) and Gundersen (2020).

were the first major areas in the United States where shale technology was deployed on a large scale. The Permian play across Texas and New Mexico have been rapidly developed during the last years of our sample period, but development came along earlier on the Texan side of the border. We thus focus our attention to the part of the Permian shale play that lie in Texas.

Estimation results for shale play level data

Subsample	Anadarko (OK)	Mississippian (OK)	Bakken (ND)	Eagle Ford (TX)	Permian (TX)
	$\ln(q_{it})$	$\ln(q_{it})$	$\ln(q_{it})$	$\ln(q_{it})$	$\ln(q_{it})$
η_{oil}	-0.10*	-0.16**	-0.12***	-0.06	0.01
	(0.06)	(0.06)	(0.06)	(0.04)	(0.07)
η_F	0.81**	1.62***	0.93***	0.33	1.09***
	(0.39)	(0.38)	(0.27)	(0.41)	(0.41)
η_{gas}	0.03	-0.03	-0.07**	-0.02	-0.06
	(0.04)	(0.04)	(0.03)	(0.04)	(0.04)
$\eta_{oil} + \eta_F$	0.70**	1.46***	0.81***	0.28	1.07***
	(0.35)	(0.34)	(0.24)	(0.37)	(0.37)
Macro controls	Yes	Yes	Yes	Yes	Yes
Well FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Well Age FE	Spline ³	Spline ³	Spline ³	Spline ³	Spline ³
First observation	2006:M09	2008:M12	2005:M02	2006:M11	2005:M02
Last observation	2017:M12	2017:M12	2017:M12	2017:M12	2017:M12
N	1,390	1,658	12,195	14,869	3,089
$N \times T$	74,474	74,489	646,096	655,691	104,092
\bar{R}^2	0.79	0.79	0.62	0.72	0.74
Clustering	Well-Time	Well-Time	Well-Time	Well-Time	Well-Time
Num. clusters	136	108	155	134	155

Table 7. Estimation results for each individual U.S. shale plays with data on log-level. A shale play is a geological formation where unconventional oil reserves are prevalent. Parameters η_{oil} and η_{gas} are the coefficients on the natural log of WTI and Henry Hub spot prices. η_F is the coefficient on natural log of the spot-futures spread. $(\eta_{oil} + \eta_F)$ is estimated by an auxiliary model. All wells are shale wells. Bakken and Permian shale plays cross state lines and we exclude Montana and New Mexico from the estimations, respectively. This to eliminate any possible confounding factors the different jurisdictions can cause.

The results for the five shale plays are reported in Table 7. For the four shale plays, Mississippian, Anadarko, Bakken and Permian, we find price responses that are statistically significant and that exceeds the average baseline response of all shale producers (0.62). However, the results for Eagle Ford differ as the prices response is not statistically significant and lower than the average baseline response. A closer scrutiny of the Eagle

Ford production data suggests that this may be due to a major structural break in 2014, after which volatility increased and production declined. While we are uncertain about the exact reason for the break in data, we find that if we redo the analysis by including an interaction-dummy that is equal to one for observations starting in 2015, the price responsiveness of Eagle Ford are in line with those of Bakken. We therefore infer that wells that are located within *shale plays* respond stronger on average than wells located outside shale plays.

4 Conclusion

We investigate the price-responsiveness of U.S. oil producers. With a novel well-level dataset covering ten of the largest oil producing states, we construct a rich panel dataset and estimate a fixed effects model. Because shale wells need to be fracked in order to start production, it introduces a new margin for the producers to exploit that conventional producers do not have. In particular, shale producers is confronted with an option to postpone production and thus are able to better time their revenue stream to more favourable market conditions.

We find shale wells to responds strongly positive to expected increases in the price of oil as measured by the spot-futures spread. In particular, across all ten states in our sample, the estimated response is 0.62. Examining the geography dimension of the cross-section we find that there is some heterogeneity across states, but except for three states (California, New Mexico and Wyoming) the responses are positive and statistically significant. The three states with insignificant (albeit positive) response are either topographically separated from the rest of the U.S. crude oil market or have insufficient pipeline infrastructure to bring potential output to market. We further show that the responsiveness depends on the level of well output level, and the strongest price-responses are found for wells operated by the largest firms and wells that are sufficiently spaced apart from their closest neighbouring well.

Our empirical results calls for new models that can account for a growing share of shale oil in the U.S., the inherent flexibility of shale extraction technology in production and the role of shale oil in transmitting oil price shocks to the U.S. economy. Our results can also serve to reconcile some of the opposing conclusions in the literature when it comes to how one should analyze the role of oil in the macroeconomy. In particular, oil price-macro models have often assumed aggregate oil production to be price inelastic in the short run when identifying oil market shocks. 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 identification

strategies for oil market macro models that relax the assumption of a zero short-run oil supply elasticity.

References

- ADRIAN, T., N. BOYARCHENKO, AND D. GIANNONE (2019): “Vulnerable growth,” *American Economic Review*, 109(4), 1263–89.
- ALLCOTT, H., AND D. KENISTON (2017): “Dutch disease or agglomeration? The local economic effects of natural resource booms in modern America,” *The Review of Economic Studies*, 85(2), 695–731.
- ANDERSON, S. T., R. KELLOGG, AND S. W. SALANT (2018): “Hotelling under pressure,” *Journal of Political Economy*, 126(3), 984–1026.
- ASHEIM, G. B., T. FÆHN, K. NYBORG, M. GREAKER, C. HAGEM, B. HARSTAD, M. O. HOEL, D. LUND, AND K. E. ROSENDAHL (2019): “The case for a supply-side climate treaty,” *Science*, 365(6451), 325–327.
- BAUMEISTER, C., AND J. D. HAMILTON (2019): “Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks,” *American Economic Review*, 109(5), 1873–1910.
- BJØRNLAND, H. C., F. M. NORDVIK, AND M. ROHRER (2021): “Supply Flexibility in the Shale Patch: Evidence from North Dakota,” *Journal of Applied Econometrics*, 36, 273–292.
- BJØRNLAND, H. C., AND J. ZHULANOVA (2018): “The shale oil boom and the US economy: Spillovers and time-varying effects,” Manuscript: BI Norwegian Business School.
- BORNSTEIN, G., P. KRUSELL, AND S. REBELO (2021): “A World Equilibrium Model of the Oil Market,” *Review of Economic Studies*, Forthcoming.
- CALDARA, D., M. CAVALLO, AND M. IACOVIELLO (2019): “Oil price elasticities and oil price fluctuations,” *Journal of Monetary Economics*, 103, 1–20.
- DAHL, C., AND M. YUCEL (1991): “Testing Alternative Hypotheses of Oil Producer Behavior,” *The Energy Journal*, Volume 12(Number 4), 117–138.
- FÆHN, T., C. HAGEM, L. LINDHOLT, S. MÆLAND, AND K. E. ROSENDAHL (2017): “Climate policies in a fossil fuel producing country—demand versus supply side policies,” *The Energy Journal*, 38(1).
- FIRPO, S., N. M. FORTIN, AND T. LEMIEUX (2009): “Unconditional quantile regressions,” *Econometrica*, 77(3), 953–973.

- GUNDERSEN, T. S. (2020): “The impact of US supply shocks on the global oil price,” *The Energy Journal*, 41(1).
- HAMILTON, J. D. (2009): “Understanding Crude Oil Prices,” *The Energy Journal*, 30(2), 179–206.
- KILIAN, L. (2009): “Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market,” *American Economic Review*, 99(3), 1053–1069.
- (2016): “The impact of the shale oil revolution on US oil and gasoline prices,” *Review of Environmental Economics and Policy*, 10(2), 185–205.
- 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.
- (2014): “The role of inventories and speculative trading in the global market for crude oil,” *Journal of Applied Econometrics*, 29(3), 454–478.
- NEWELL, R. G., AND B. C. PREST (2019): “The Unconventional Oil Supply Boom: Aggregate Price Response from Microdata,” *The Energy Journal*, 40(3).
- 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.
- RIOS-AVILA, F. (2019): “Recentered influence functions in Stata: Methods for analyzing the determinants of poverty and inequality,” *Levy Economics Institute, Working Paper*, 927.
- SMITH, J. L. (2009): “World Oil: Market or Mayhem?,” *Journal of Economic Perspectives*, 23(3).

Centre for Applied Macroeconomics and Commodity Prices (CAMP)

will bring together economists working on applied macroeconomic issues, with special emphasis on petroleum economics.

BI Norwegian Business School
Centre for Applied Macro - Petroleum economics (CAMP)
N-0442 Oslo

www.bi.no/camp