CENTRE FOR APPLIED MACROECONOMICS AND COMMODITY PRICES (CAMP)

CAMP Working Paper Series
No 8/2019

Supply flexibility in the shale patch: Facts, no fiction

Hilde C. Bjørnland

BI

Supply flexibility in the shale patch: Facts, no fiction*

Hilde C. Bjørnland†

November 18, 2019

Abstract

In two recent papers, Kilian and Zhou (2019) and Kilian (2019) have criticized Bjørnland, Nordvik, and Rohrer (2017), arguing that our finding of a large price elasticity of output for shale producers is not credible. We welcome a discussion of our methods and findings, but the criticisms made in these two papers are inaccurate and mischaracterize our analysis and results. In this note I address the criticism that has been made, arguing that our findings support the notion that the degree of output flexibility is dependent on the production technology in question. Furthermore, I argue that knowledge that shale producers are more price elastic than conventional oil producers could have far reaching implications for the industry, for macroeconomic outcomes, and for policy analysis.

*The author would like to thank Knut Are Aastveit, Jamie Cross, Thomas S. Gundersen, Even C. Hvinden, Frode M. Nordvik, Maximilian Rohrer and Leif Anders Thorsrud for useful comments and fruitful discussions. The usual disclaimers apply. The views expressed in this paper are those of the author and do not necessarily reflect those of Norges Bank.

†Centre for Applied Macroeconomics and Commodity prices (CAMP) - BI Norwegian Business School and Norges Bank. Email: hilde.c.bjornland@bi.no

1 Introduction

Are shale oil producers responding to current and future oil prices by changing oil production and completion of new oil wells? Yes, according to a recent panel data study by Bjørnland, Nordvik, and Rohrer (2017) (hereafter BNR), that investigates the response of both shale and conventional crude oil wells to spot and expected future oil prices. In particular, constructing a novel well-level/firm-level monthly production data set from North Dakota, we find the supply elasticity of shale wells to be significantly positive and in the range of 0.3-0.9, depending on wells and firms characteristics. We find no such responses for conventional wells, and interpret the supply pattern of shale oil wells to be consistent with Hotelling (1931)'s model of optimal extraction, with flow production as the relevant control variable.

These are new findings in the literature, and have important implications for how one should analyze the role of oil supply shocks for the macroeconomy. Until recently, oil price-macro models have often assumed the short run price elasticity of aggregate oil production to be zero, or at least, small, when identifying oil market shocks, see for instance Kilian (2009) and Kilian and Murphy (2012), and a series of other papers building on the seminal paper of Kilian (2009). While this may be consistent with the behaviour of *conventional* producers, c.f. Anderson, Kellogg, and Salant (2018), our results for *shale* producers support exploring alternative identification schemes that relax the assumption of a zero short-run oil supply elasticity, such as the Bayesian approach recently developed in Baumeister and Hamilton (2019a). Doing so, they find oil supply disruptions to be a bigger factor in historical oil price movements than implied by the earlier estimates in studies such as Kilian (2009) and Kilian and Murphy (2012).

The paper of Baumeister and Hamilton (2019a) has been criticised by Kilian and Zhou (2019), stating, among others, that the results are driven by the imposition of a highly unrealistic prior for the impact price elasticity of global oil supply: a supply elasticity which is supported by the findings in BNR, among others. In particular, Kilian and Zhou (2019) dismiss the results of BNR, based on two grounds: (i) our paper is only based on supply responses in North Dakota and (ii) the results are based on the prices of contracts for future delivery, hereafter futures. Ignoring the coefficient on futures prices, they claim that we find an insignificant (or even negative) supply elasticity, in line with what Kilian (2009) and Kilian and Murphy (2012) assume. The critique put forward has been thoroughly addressed by Baumeister and Hamilton (2019b).[1] Then, in a follow up paper, Kilian (2019) has added a series of additional concerns related to the papers

[1] Kilian and Zhou (2019) raise additional concerns with regard to the approach of Baumeister and Hamilton (2019a), unrelated to this paper. In their response, Baumeister and Hamilton (2019b) have shown that the criticism is inaccurate, both in broad substance and in specific details. In particular, they argue that

by Baumeister and Hamilton (2019a,b), and the literature on the supply elasticity. The paper, which is entitled: "Facts and Fiction in Oil Market Modeling", concludes that the result of a large oil supply elasticity found in BNR is not credible. Instead, Kilian (2019) refers to the results of a recent paper by Newell and Prest (2019), which analyses the elasticity at distinct phases of well development, i.e., drilling, completion and production for shale producers in five different states, but find the short run supply elasticity to be zero.

In this paper I discuss the criticism put forward by Kilian and Zhou (2019) and Kilian (2019) with regard to the supply elasticity literature. In particular, I describe how these two papers mischaracterise the model and results found in BNR, and in so doing, downplay the importance and implication of shale oil producer's behaviour. BNR have shown that the production of horizontally drilled shale oil differs from conventional operations. Furthermore, we argue that the option of executing so-called drilled but uncompleted (DUC) wells suggest a plausible mechanism through which shale oil firms may respond to price developments by increasing production. Our empirical results therefore support the notion that the degree of output flexibility is dependent on the production technology in question. This is a new finding in the literature, that could have far reaching implications for the industry, for macroeconomic outcomes, and for policy analysis.

This paper is organized as follows. Section 2 reviews the key results of BNR, while in Section 3 I address the main criticisms put forward by Kilian and Zhou (2019) and Kilian (2019) with regard to the oil supply literature in general, and BNR in particular. Furthermore I give some suggestions for why the results of BNR differ from those of Newell and Prest (2019). I conclude in Section 4 by also discussing the wider implications of the findings.

2 Key results of Bjørnland, Nordvik, and Rohrer (2017) and relations to the literature

In a competitive market the extraction technology of even a subset of oil producers may substantially influence the aggregate elasticity of oil supply. The emergence of the shale oil boom has not only doubled U.S. oil output, but starkly altered the manner in which it is produced. Knowledge of producer behaviour and supply elasticity in unconventional (shale) oil pools is therefore important, but until recently, it has been lacking.

Why should shale producers behave differently from conventional oil producers? As shown in BNR, a shale well may be stimulated with water and chemicals many times

Kilian and Zhou (2019) mischaracterize the literature on supply elasticity, and in particular the analysis and results of BNR and Caldara et al. (2019), see Baumeister and Hamilton (2019b) for details.

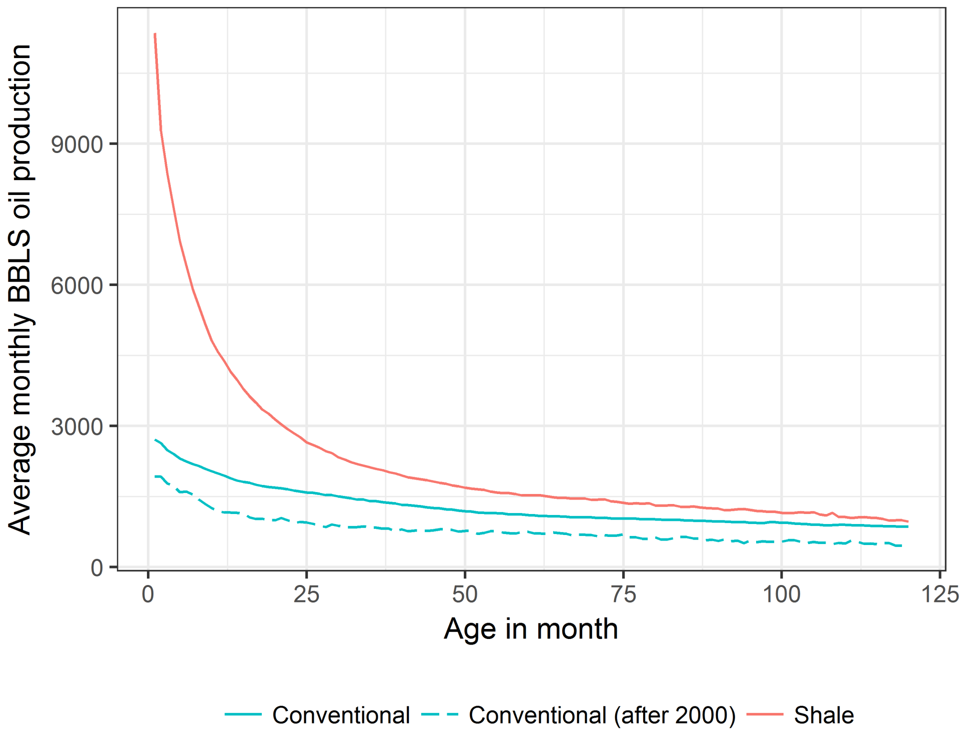

Figure 1. Production profile of horizontal and vertical drilled wells

Note: The figure is identical to Figure 2 in Bjørnland et al. (2017). It reports the average monthly production of BBLS oil during the first 120 months after spudding. We exclude the first month of production as well typically do not operate every day of this month. Wells are separated in conventional/vertical wells, indicated by the blue color, and unconventional/horizontal/shale wells, indicated by the red color. The solid line cover wells spudded between February 1952 and June 2017, while the dashed line only considers conventional wells spudded after January 2000. See Bjørnland et al. (2017) for details on data.

during its lifetime. This gives the well operator a certain degree of freedom to choose the timing of the fracking operations relative to a conventional well which is naturally flowing. Furthermore, shale wells have a very front-loaded production profile relative to conventional wells. In particular, in the first two years of the lifetime of a shale well, production declines by a monthly rate twice as high as the average decline rate for conventional wells, see Figure 1, which is taken from BNR. This production front-loading increases the incentive for shale producers to optimise the timing of well completion and production.

In BNR, we formally analyse to what extent shale oil producers respond to prices by changing oil production at the well level, i.e., supply elasticity, and the completion of new wells. We analyse this using both well-level data and firm-level data. In the main specification, we include lagged production changes, various price signals and a set of controls. Importantly, we include both changes in the spot prices and changes in spot-futures spreads. The futures market for crude oil provides intertemporal price signals, that producers also consider closely. In particular, if spot prices increase relative to the future, we would expect the well owners to maximize profit by increasing production today, and vice versa if the spread is decreasing. We expect no relationship for conventional wells, that are less flexible than shale. Our well-level monthly production data set covers more than 16,000 crude oil wells in North Dakota, including both shale and conventional wells. We find the supply elasticity of shale wells to be significantly positive and in the range of 0.3-0.9, depending on wells and firms characteristics. In particular, the largest response is found for younger wells (wells younger than three years) and for large firms (top 99 firms with highest production volumes).

Since we first put our paper on our webpage, corroborating results are found in Bornstein, Krusell, and Rebelo (2018) using annual data to study production of shale producers. Arguably, it appears that firms exploit the inherent flexibility of shale technology to allocate production volumes intertemporally. We interpret the results for shale oil producers to be in line with Hotelling (1931)'s model of optimal exhaustible resource extraction. Reserves are an inventory, and the decision to produce is an intertemporal choice of when to draw down below-ground inventory. For producers to behave in line with the Hotelling (1931) theory, they must be able to reallocate extraction across different periods.

Our results are in contrast with recent literature on conventional oil. Notably, Anderson et al. (2018) have shown that the price elasticity of conventional oil producers in Texas is statistically indistinguishable from zero. They present a strong argument that the relevant conventional control variable for oil firms is exploration investment, not flow production. However, the production of horizontally drilled shale oil differs from conventional operations. In particular, the option of executing so-called drilled but uncompleted (DUC) wells suggest a plausible mechanism through which shale oil firms may respond to price developments by increasing production. Our empirical results support the notion that degree of output flexibility is indeed contingent on the production technology in question.

These are new results in the literature, that have potentially important implications. As shale producers grow in size and importance, we should expect to see a stabilizing effect on oil prices. Furthermore, the results have important implications for how one should analyze the role of oil supply shocks in the macroeconomy, as the elasticities of supply and demand are crucial for determining the relative importance of oil supply vs oil demand shocks for oil price fluctuations and macroeconomic behaviour, see Baumeister and Hamilton (2019a) and Caldara et al. (2019). Finally, knowledge about shale behaviour may also have important policy implications. For instance, when policy makers are designing tax policies that affect the petroleum industry, taking into account that shale producers can adjust differently to price-sensitive news than what conventional producers can do, could be important. I will discuss this in more detail in Section 4, after having responded to the criticism put forward by Kilian and Zhou (2019) and Kilian (2019) in Section 3.

3 Discussion of Kilian and Zhou (2019) and Kilian (2019)

The recent paper by Kilian (2019) argues, among others, that the results of a large supply elasticity reported in BNR are not credible. In so doing, the paper reiterates some of the criticism put forward by Kilian and Zhou (2019). Below I address each of their concerns.

First, both Kilian and Zhou (2019) and Kilian (2019) argue that including a futures spread to capture expectations has no purpose. For this reason, they ignore the coefficient on the spread reported in BNR, focusing only at the coefficient on spot price, which is negative. This would suggest that the supply curve is downward sloping. However, this coefficient can not be interpreted in isolation. In so doing, they are effectively assuming that changes in production from a well in response to an increase in the current price is exactly matched by an increase in the futures price. Hence, the spread is zero. Our results suggests otherwise. The coefficient on the spread is significantly positive.

Why should the futures spread matter for the supply elasticity? U.S. shale producers can lock in prices for their production months into the future, and then sell at opportune moments. Importantly, the futures market for crude oil provides intertemporal price signals that producers also consider closely. For instance, if spot prices increase relative to the future, one would expect the shale well owners to maximize profit by increasing production today, and vice versa if the spread is decreasing.[2] NYMEX futures are traded liquidly at the time horizons considered here, and with many risk-neutral traders, the futures price should be a reasonable approximation of the expected future spot price. This is also consistent with how oil well operators are believed to use the futures market to make price projections, see for instance Newell and Prest (2019), footnote 10: " *Using 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.*" As mentioned in the introduction, Kilian (2019) has confidence in the paper by Newell and Prest (2019), that analyze shale oil behaviour in five different U.S. states. Interestingly, they also use future prices, which is not commented on in Kilian (2019). In fact, they use the average of the next 12 months of futures prices for the WTI and Henry Hub prices instead of the spot prices in the parts that analyse drilling

[2] Central banks and other policy agencies typically use future oil prices while making forecasts. In particular, the slope of the yield curves formed with these prices provide information on the direction of spot prices, see Reeve and Vigfusson (2011) for an analysis of the usefulness of future prices.

and completion of the wells, (see page 5), while they use only the spot prices when they analyse well production. So while Kilian (2019) criticises the use of futures prices in BNR, it goes unnoticed that Newell and Prest (2019) base all their models with positive results on futures prices. Still, even if one did not believe that future prices should be relevant for the output choice, ignoring the coefficient once it is estimated, like Kilian and Zhou (2019) and Kilian (2019) do, makes no sense. Furthermore, if we instead include spot prices and futures prices separately (which is also consistent with what Anderson et al. (2018) do) our results are robust, see Table A.5. in BNR. Finally, since oil producers respond significantly to futures prices, and the spot and futures prices are correlated, excluding the futures price may result in a biased estimate of the spot price elasticity.

Second, Kilian (2019) argues that there is no compelling reason for including oil futures prices as a proxy for oil price expectations when estimating the oil supply elasticity, due to storage demand. He argues that if there are exogenous shifts in expectations about future oil prices, this will cause a change in storage demand, which in turn shifts the spot price. Hence, the model already captures this effect by including changes in the spot price and allows producers to respond to this type of shock. While I agree that storage definitely has a role to play in macroeconomic models, as price moves with changes in inventories (above ground) driven by shifts in inventory demand, the argument is not relevant at the well level. Individual well owners or firms do not hold storage above ground. Whatever they store is underground, and exactly for that reason including future prices has a purpose. At the monthly frequency, there is also strong evidence, that this information is relevant.

Third, and related to the above. While critical to including oil futures prices in the analysis, Kilian (2019) acknowledges that oil suppliers could respond to higher *expected* oil prices by storing oil *below* the ground. Shale producers have this option, for technological reasons, as they may drill, but not frack a well in anticipation of rising prices. Recall that these drilled, but not yet completed wells are known as DUCs. We show in an extensive analysis that DUCs do indeed respond strongly to price signals, suggesting that production could also be a choice variable. Still, Kilian (2019) dismisses our results on the DUCs, arguing that operators can not execute the DUC option in response to higher oil prices, as "*it still takes between four and twelve weeks to complete the well*". This is not correct. First, to efficiently schedule well completions, producers maintain a reasonable number of DUCs for operational flexibility and/or economic reasons, see for instance EIA (2019). The length of time that an oil well has been drilled, but remains uncompleted, can vary a lot. Some oil wells are completed immediately (within days) after drilling, but other wells remain drilled but uncompleted for several months or even years. This flexibility is in particular notable for shale wells. As shown in Figure 3 in BNR, the distribution of time between drilling/spudding and completion of a well is more dispersed for shale wells than for conventional. The larger variation in completion time supports the notion that deciding to complete a well is a choice variable, which motivates why we study supply elasticity also on the extensive margin. In contrast to what Kilian (2019) concludes, our reported findings for the DUCs are not only important, but they also suggest that the results we have found for production are credible.

Fourth, Kilian (2019) questions the differencing of the spread measure. BNR use the difference in growth rates between the spot price and the futures price for a contract with delivery at time $t + j$ to measure the expected change in market conditions between the spot market for crude oil and the future market of oil j months ahead. Our focus is on the change in direction. If this measure is positive, the spot market looks relatively more favorable measured against the futures market than before. If it is negative, the future market j months ahead has become relatively less favorable than before. Our choice of differencing the spread can also be supported from a time series perspective: neither the oil price nor the spread are stationary, supporting that we should difference both these series. Furthermore, and as mentioned above, we also analyse robustness to this specification by including the difference in the spot prices and the difference in futures prices separately, and get similar result, again see Table A.5. in BNR. Our specification is also consistent with what has been chosen in many other studies, including Anderson et al. (2018) for conventional oil, that are differencing all their variables, including the futures. In contrast, Newell and Prest (2019) do not find oil flow supply to be price elastic even for shale producers. However, they include oil prices in levels.[3] Furthermore, in addition to oil prices, they also include gas prices in the same estimation, potentially correlated with oil prices. In our data set, neither the oil price, nor the spread are stationary, which is an added reason for why we have differenced the data.

Fifth, Kilian (2019) argues that the estimated coefficients are not valid because of strong collinearity. He argues that he can rewrite the specification and will then find a correlation between the variables to be 98%. It is unclear as to how he found that correlation, as the correlation between the changes in the spot price and changes in the spot futures spread in our data is between 40% and 65%, depending on specification and sample. Furthermore, we perform extensive robustness tests, showing that the results hold across different specifications, such as changing the sample, the maturity of the futures contract, or using well-level or firm-level data. The only changes we find is when we separate for firms, i.e., we find that the largest response is for younger wells (wells younger than three years) and for large firms (top 99 firms with highest production volumes).

[3] To be precise, Newell and Prest (2019) use different specifications throughout the paper. They estimate the relationship between drilling and prices in first differences "to make the revenue and drilling series stationary" (see footnote 14), while for production, prices are specified in levels (see page 16).

Sixth, Kilian (2019) argues that since our results are not consistent with Anderson et al. (2018), what we find is difficult to reconcile with economic theory (see footnote 14). The problem with this argument is that Anderson et al. (2018) is about conventional oil production. We analyze unconventional oil extraction with a different production technology. Anderson et al. (2018) refute the original Hotelling model because it does not fit with conventional producers in Texas, and therefore write a new augmented Hotelling model, where producers do not vary flow supply in response to price signals. However, we show that the production from horizontally drilled shale wells differs from conventional operations. In particular, the option of executing DUC wells suggests a plausible mechanism through which shale oil firms may respond to price developments by increasing production. Our empirical results thus support the notion that the degree of output flexibility is contingent on the production technology in question. Hence, we do not need an augmented model, like the Anderson et al. (2018) model. Shale well technology is more in line with the original Hotelling (1931) theory. See also the paper by Bornstein et al. (2018) that finds corroborating results (using annual data).

On a final note, and as mentioned above, Kilian (2019) has confidence in the results of Newell and Prest (2019) that do not find oil flow supply to be price elastic. There are, however, a few additional differences between this study and BNR. First, Newell and Prest (2019) do not explore well and firm characteristics, as BNR do. By constructing a rich panel data set, we can eliminate any potential aggregation bias over well production rates when estimating the empirical model. Second, and as mentioned above, only for completion and investment do they include future price signals, whereas BNR include this in all specifications. In ongoing work, Aastveit, Bjørnland, and Gundersen (2019) estimate monthly supply elasticity for all major oil producing US states using novel micro data from Rystad Energy. This analysis confirms the conclusion drawn in BNR for North Dakota, but also show that most shale producers in the U.S. respond positively and significantly to oil price signals, with the size depending somewhat on the density of the wells. What is more, we can encompass the results in Anderson et al. (2018) for conventional oil production in Texas, but add new information by showing that shale producers behave differently (which were not part of their original sample). With regard to the results reported in Newell and Prest (2019), we show how the results will depend on the appropriate data transformation, see Aastveit et al. (2019).

4 Concluding remarks and wider implications

The main takeaway from Bjørnland, Nordvik, and Rohrer (2017), is that shale oil producers are responding to current and future oil prices by changing oil production and completion of new oil wells. In particular, we find the supply elasticity of shale wells to be significantly positive: the size depending somewhat on wells and firms characteristics.

While these results are for North Dakota only, the results have far reaching implications. By 2017, oil production from shale deposits accounted for half of US crude oil output. Furthermore, the use of hydraulic fracturing technology is spreading to other oil-producing countries, potentially making unconventional oil a much larger share of total production than it is today. This could imply a stabilizing effect on oil prices, as has recently also been emphasized in Bornstein et al. (2018). Furthermore, the results of BNR could also have far reaching implications for policy. In particular, the inherent flexibility could also be a factor to consider when designing tax policies that affect the petroleum industries. Acknowledging that shale producers can adjust differently to price-sensitive news than conventional producers, should be important.

As already emphasized, our results run contrary to how many oil macro models, such as Kilian and Murphy (2012) and Kilian and Murphy (2014), identify oil price shocks through imposing a very small price elasticity of supply. In so doing, they can not account for the flexibility of shale producers. Instead our results support exploring identification schemes like those of Baumeister and Hamilton (2019a), that relax the assumption of an extremely small short-run oil supply elasticity, or Caldara et al. (2019), that combine a narrative analysis of episodes of large drops in oil production with country-level instrumental variable regressions. Doing so, both find oil supply disruptions to be a bigger factor in historical oil price movements than implied by previous studies.

Finally, although not directly relevant for the results here, I note that the shale oil revolution has also changed other aspects of the workings of the U.S. economy, including a wide range of outcomes such as local labor market dynamics, investment, and real activity. In particular, a growing literature using panel data studies have recently shown that the shale oil boom has benefited non-oil employment and wages at the local level in many resource abundant U.S. states, c.f. Allcott and Keniston (2017). What's more, there may be potential spillovers to other parts of the US economy through technological advancement and learning by doing, emphasized by Bjørnland and Zhulanova (2018).[4]

To conclude, the empirical results of BNR support the notion that the degree of output flexibility is indeed contingent on the production technology in question. This calls for new models that can account for a growing share of shale in the U.S., shale's inherent flexibility in production, and the role of shale oil in transmitting oil price shocks to the U.S. economy.

[4]See also Bjørnland and Thorsrud (2016) and Bjørnland, Thorsrud, and Torvik (2019) that analyse spillovers of oil and gas in Australia and Norway.

References

Aastveit, K. A., H. C. Bjørnland, and T. S. Gundersen (2019). Oil Supply Elasticities Across the U.S.: Micro Evidence and Macro Implications. Manuscript, BI Norwegian Business School.

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.

Baumeister, C. and J. D. Hamilton (2019a). Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks. *American Economic Review 109*(5), 1873–1910.

Baumeister, C. and J. D. Hamilton (2019b). Structural interpretation of vector autoregressions with incomplete identification: Setting the record straight. Manuscript, UC San Diego.

Bjørnland, H. C., F. M. Nordvik, and M. Rohrer (2017). Supply flexibility in the shale patch: Evidence from North Dakota. CAMP Working Paper 2/2017 - Manuscript revised May 2019, BI Norwegian Business School.

Bjørnland, H. C. and L. A. Thorsrud (2016). Boom or gloom? Examining the Dutch disease in two-speed economies. *The Economic Journal 126*(598), 2219–2256.

Bjørnland, H. C., L. A. Thorsrud, and R. Torvik (2019). Dutch disease dynamics reconsidered. *European Economic Review 119*, 411–433.

Bjørnland, H. C. and J. Zhulanova (2018). The shale oil boom and the US economy: Spillovers and time-varying effects. CAMP Working Paper 8/2018, BI Norwegian Business School.

Bornstein, G., P. Krusell, and S. Rebelo (2018). Lags, Costs, and Shocks: An Equilibrium Model of the Oil Industry. Working Paper 23423, National Bureau of Economic Research.

Caldara, D., M. Cavallo, and M. M. Iacoviello (2019). Oil price elasticities and oil price fluctuations. *Journal of Monetary Economics 103*, 1–20.

EIA (2019). Today in Energy, September 10. U.S. Energy Information Administration.

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

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

Kilian, L. (2019). Facts and Fiction in Oil Market Modeling. Manuscript, September, Federal Reserve Bank of Dallas.

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.

Kilian, L. and D. P. Murphy (2014). The role of inventories and speculative trading in the global market for crude oil. *Journal of Applied Econometrics 29*(3), 454–478.

Kilian, L. and X. Zhou (2019). Oil supply shock redux? Manuscript, July, University of Michigan.

Newell, R. G. and B. C. Prest (2019). The Unconventional Oil Supply Boom: Aggregate Price Response from Microdata. *The Energy Journal 40*(3).

Reeve, T. A. and R. J. Vigfusson (2011). Evaluating the forecasting performance of commodity futures prices. FRB International Finance Discussion Paper 1025.

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

CAMP Working Paper Series
ISSN: 1892-2198