DEPARTMENT OF ECONOMICS
OxCarre
Oxford Centre for the Analysis of Resource Rich Economies

Manor Road Building, Manor Road, Oxford OX1 3UQ
Tel: +44(0)1865 281281 Fax: +44(0)1865 271094
oxcarre@economics.ox.ac.uk www.oxcarre.ox.ac.uk

OxCarre Research Paper 152

The Relationship Between Oil Price and Costs in the Oil and Gas Industry

Gerhard Toews

OxCarre

&

Alexander Naumov

Economics Team, BP

Direct tel: +44(0) 1865 281281 E-mail: oxcarre@economics.ox.ac.uk

The Relationship Between Oil Price and Costs in the Oil and Gas Industry *

Gerhard Toews[a†] and Alexander Naumov[b‡]

[a] *Oxford Centre for the Analysis of Resource Rich Economies, Department of Economics, University of Oxford*

[b] *Economics Team BP*

January 2015

Abstract:
We propose a simple structural model of the upstream sector in the oil and gas industry to study the determinants of costs with a focus on its relationship with the price of oil. We use the real oil price, data on global drilling activity and costs of drilling to estimate a three-dimensional VAR model. We use short run restrictions to decompose the variation in the data into three structural shocks. We estimate the dynamic effects of these shocks on drilling activity, costs of drilling and the real price of oil. Our main results are that (i) a 1% increase (decrease) in the oil price increases (decreases) global drilling activity by 1% and costs of drilling by 0.5% with a lag of a year; and (ii) shocks to drilling activity and costs of drilling do *not* affect the price of oil permanently.

Keywords: Natural Resource Extraction, Crude Oil Price, Upstream Cost
JEL classification: Q31

*We would like to thank Paul Collier, Martina Kirchberger, Zeyyad Mandelinici, Tony Venables and the participants of the OxCarre Brown Bag Seminar, SMYE 2014, Gorman seminar at the University of Oxford and BP Economics Team for their useful comments and suggestions.

†Correspondence: Oxford Centre for the Analysis of Resource Rich Economies, Department of Economics, Manor Road Building, Oxford OX1 3UQ, UK; Email: gerhard.toews@economics.ox.ac.uk

‡Correspondence: BP, 1 St. James's Square, London, SW1Y 4PD; Email: alexander.naumov@bp.com

1 Introduction

The economic profession is still struggling to fully understand the crude oil price and the shocks driving it (Hamilton, 2008; Kilian, 2009; Venables, 2011; Anderson, Kellogg, and Salant, 2014). The goal of this paper is to contribute to understanding the determinants of drilling costs and the relationship between drilling costs and the real price of oil. We use data on individual wells from Wood MacKenzie on drilling activity and costs of drilling. Focusing on the information provided for the biggest 25 oil and gas companies we construct two quarterly time series capturing (i) the total number of exploration wells drilled by these companies and (ii) the average cost of drilling these wells. We use the constructed time series in combination with the real price of oil to estimate a three-dimensional structural VAR model.

The proposed structural model of the **upstream sector** allows us to decompose the variation in the reduced form errors of the estimated VAR into three structural shocks. To identify the shocks we assume a recursive structure similar to Kilian (2009). The first structural shock is an *oil price demand shock* which is defined as an unpredictable innovation to the oil price. To identify the shock we assume that the oil price is predetermined to drilling activity and costs of drilling. This assumption is plausible because it takes more than three months, and typically more than a year, following the drilling of an exploration well before the global supply of oil and gas can be affected. Moreover, expectations of the future oil price are unlikely to be formed contemporaneously because it takes more than a year before reasonable estimates of reserves and drilling costs are available.

The remaining variation in the errors - after accounting for the variation in the oil price - is decomposed into a *supply shock* and a *drilling specific demand shock*.

Supply shocks in the upstream sector are driven by productivity changes of firms providing drilling services, e.g. Halliburton, Schlumberger, Baker Hughes. Drilling specific demand shocks arise from productivity changes of firms demanding drilling services, e.g. BP, Shell, Chevron. The latter is referred to as 'drilling specific' to emphasize that the shock originates in the upstream sector and is *not* driven by fluctuations in the oil price. We disentangle the two shocks by assuming that the number of exploration wells which can be drilled in the short run is fixed (vertical supply curve). Thus, only supply shocks can affect the number of wells drilled and costs of drilling contemporaneously. We justify this assumption by arguing that the exploration of new territories is associated with the necessity to utilize and transport rigs, and contract, maybe even educate, a new crew to operate the rig in a new environment.

We have two main results. First, following a 1% increase (decrease) in the oil price drilling activity increases (decreases) by 1% and costs of drilling increases by 0.5% with a lag of a year. Thus, we confirm the theoretically expected fact that the extraction of natural resources is an increasing costs industry. Second, shocks originating in the upstream sector (supply shocks and drilling specific demand shocks) do *not* have a permanent effect on the price of oil. Hence, shocks in the upstream sector do not appear to be important in determining the price of oil in the long run. Partly, this result may be explained by the endogenous expansion of crude oil production as a reaction to an increase in drilling costs (Kilian, 2009).

Our paper is most closely related in methodology to Kilian (2009) who uses a three-dimensional structural VAR and short-run restrictions to understand the different nature of shocks driving the oil price. As opposed to our approach, he models the world market for crude oil rather than the upstream sector of the oil and gas industry. To do this he uses world oil production to capture supply shocks. To model demand shocks for industrial commodities, he uses dry cargo single voyage ocean freight rates. After accounting for shifts in demand and supply, he argues that the residual variation in the oil price is driven by precautionary demand shocks rather than by supply shocks as has been previously believed. Our paper is also related to Anderson, Kellogg, and Salant (2014). Using data from Texas on drilling activity and rig rents, they present evidence that drilling activity and drilling costs significantly respond to a change in the oil price. On the other hand, they are not able to find any significant relationship of oil price changes and the contemporaneous extraction of oil. They use these results to motivate a theoretical model in which drilling activity is at the centre of deriving an optimal extraction path. In doing so, they are able to rationalize their empirical findings. Our work differs from their contribution in two main aspects. First, and most importantly, their main contribution is theoretical whereas our focus is on the causal identification and estimation of shocks. Second, we use data on drilling and costs of drilling covering the whole world, whereas their data is constraint to activity and rig rents in Texas.

The remainder of this paper is structured as follows. In the next section we discuss the institutional framework of the upstream sector in the oil and gas industry. Based on the discussion we derive a simple theoretical partial equilibrium framework to guide our empirical analysis. In the third section we describe the data. In the fourth section we discuss our identification and estimation strategy. In the subsequent section our results are presented before we conclude.

2 Upstream sector in the oil and gas industry

Activity in the oil and gas sector can be broadly divided into three sectors: an upstream sector, which involves locating and extracting the product located under the Earth's surface; midstream sector, which mainly involves the transportation and storage of the product; and downstream sector, which involves the processing, distribution and selling of the final product. In what follows we focus on the **upstream sector** as we are interested in the determinants of drilling costs and their interaction with the oil price.

In order to operate in an oil or gas field the company requires a licence from the owner of the mineral rights. Depending on the circumstances the owner is an individual, an institution or a state. Usually, several companies are simultaneously operating in the same field. Sometimes they are competing for the extraction of the product which results in common pool problems - most famously seen during the first oil boom in Pennsylvania (Yergin, 2011). And sometimes they collaborate by creating consortia which often involve the state as an important agent in the project. Partly as a consequence of the variety of work relationships the type of licences allocated to companies often varies and is frequently determined on a bilateral basis involving prolonged negotiations to agree on a profit sharing scheme.

Reserves are typically buried under many layers of rock and may be located onshore and offshore. Thus, drilling represents the core of activities in the upstream sector as it is essential to access the product. Types of wells drilled differ in their purpose. We broadly differentiate between two types of wells: exploratory wells[1] and produc-

[1] Throughout the paper we refer to exploratory and appraisal wells as exploratory wells as the latter can be considered to be part of the exploration process (Adelman, 1990).

tion wells. The former have the purpose of identifying new reserves. The latter are mostly drilled in known reservoirs to maintain or increase oil and gas production. We model the upstream sector by focusing on the drilling of exploratory wells for two reasons: an empirical and a theoretical. Empirically, the oil price may be considered predetermined to activity in drilling of exploration wells - but less so in case of production wells - which is an important assumption of our identification strategy and will be discussed further below. Theoretically, exploration is most closely linked to the movements in and out of new territories which results in a huge variation in geology which the drill has to overcome and drilling costs.

Once the well has been drilled the drilling rig is no longer required and can be moved. Drilling activity of any producer fluctuates with oil production outcomes from recently drilled wells and the firm's success in finding new fields in a particular location.[2] A successful exploration well drilled attracts subsequent drilling from the own or a competing company, and a dry hole does not. Thus, outsourcing drilling rigs reduces the overall capacity requirements of rigs and greatly reduces transportation and mobilization costs.[3] This explains why companies in the oil and gas sector are not vertically integrated and why drilling is typically outsourced (Kellogg, 2011).

The absence of a vertical integration in the upstream sector allows us to observe costs of drilling a well which are determined by an interplay of demand and supply

[2] In Figure V in the Appendix the success rate of drilling a well as a share of the total number of exploration and appraisal wells drilled by companies in our sample is presented. The obvious downward trend in the share of unsuccessful wells drilled suggests technological progress in identifying productive wells. Note that because we limit our sample to major oil and gas companies the 85% (80% if appraisal wells are dropped) success rate in 2012 might represent an upper bound if the major oil and gas companies are more successful in identifying productive wells.

[3] Despite the fact that producers outsource drilling they are responsible for identifying potential reserves, designing wells and writing drilling procedures because they are better informed about the specific geological conditions.

Figure I: Number of wells drilled (D) and cost of drilling (c_D) in the upstream sector

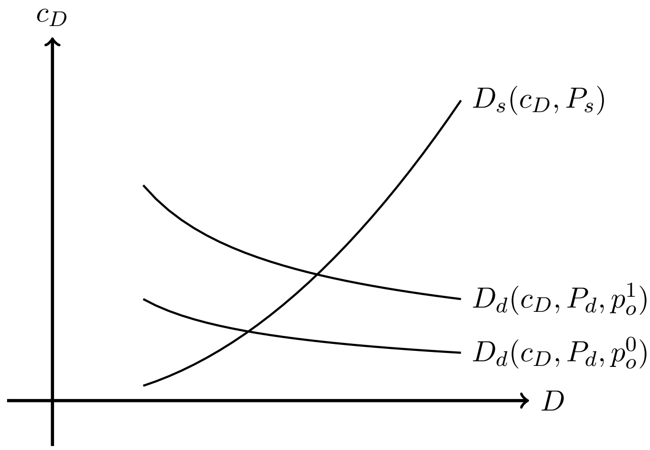

for drilling services. On the demand side oil companies (e.g. BP, Shell, Chevron, etc.) demand drilling services. On the supply side service companies (e.g. Halliburton, Schlumberger, Baker Hughes, etc.) provide drilling services. Typically, the contracting is initiated by the oil company requesting the drilling services and specifying the requirements of the hole drilled: the depth of the well to be drilled, specifics of geological layers, etc. Based on this information, the service company makes an offer by specifying the rig to be used, the crew hired and the price to be paid.

We suggest that the variation in the number of wells drilled, D, and costs of drilling, c_D, may be decomposed into three structural shocks (see Figure I). Abstracting from changes in technology and regulatory constraints, at any point in time, a small output can be drawn from lower-cost and high-quality sources. But a larger output has to be drawn from higher-cost low-quality sources. An increase in the oil price, ceteris paribus, increases the profitability of oil wells and triggers an increase in drilling activity in higher-cost low-quality areas. Thus, the first shock is an unanticipated change in the real price of oil, $\varepsilon^{p_o} = p_o^1 - p_o^0$. We think of this shock as a demand shifter in the upstream sector and we refer to this shock as an *oil price demand shock*.

Now, we relax the above assumption and allow for changes in technology (improvements in seismic technology, development of horizontal drilling, advances in deep sea offshore drilling, etc.) and regulatory constraints (expansion of area available for leases, nationalization of an oil and gas industry, increase in taxes or tougher safety regulations, etc.) to contribute to the variation in the number of wells drilled and cost of drilling. We will refer to technological and regulatory changes jointly as productivity shocks. Productivity shocks may affect firms on the supply side and on the demand side. On the supply side, ε^{P_s} is the productivity shifter affecting firms which provide drilling services. We refer to this shocks as *supply shocks*. On the demand side, ε^{P_d} is the productivity shifter affecting firms which demand drilling services. We refer to this shock as a *drilling specific demand shock*. The term 'drilling specific' is chosen to emphasize that the shift in demand is *not* driven by fluctuations in the oil price.

3 Data

In the first part of this section we discuss the source and the quality of our data. In the second part we discuss how w transform the raw data for the empirical anylsis.

The raw dataset was obtained from Wood Mackenzie, and compiled using a range of methods: (i) meetings with energy companies, annual reports and industry specific publications, (ii) state publications and information from public institutions, (iii) historical data, investor presentations and different types of media sources.

To our knowledge, this is the best source of data on drilling activities and costs in the oil and gas industry covering the whole world. Nevertheless, we are concerned about the quality of the reported data. To address these concerns we limit the data in cross section and time dimension. First, we have picked the 25 biggest oil and gas companies based on their production levels as reported by Forbes (Helman, 2013).[4] In Figure II all companies used in the empirical analysis are presented. Choosing the major companies is motivated by the idea that the information provided for and by these companies is more reliable. This is because publicly traded companies are required to be transparent. Moreover, many active interest groups watch the activity of these companies carefully (e.g. shareholders, governments, environmental activists, competitors). Restricting the analysis in this way allows us to capture the important changes in activity and costs in this sector for two reasons. In total the 25 selected companies produced more than 50% of total oil and gas production in every year of our sample. Moreover, these companies are arguably at the technological frontier and the first to engage in a new type of drilling or first to explore a new region of the world.

Second, we limit our data in the time dimension by starting our analysis in 1992 because the quality of data appears to decrease significantly going back in time (many more observations on depth and costs of drilling a well are missing before 1992). Moreover, we believe that modelling an integrated global oil market is much more realistic since the collapse of the Soviet Union in 1991.

Approximately 10% of the reported cost of drilling is missing. We replace the missing

[4] We excluded Iraq due to it's turbulent recent past. Instead we add Anadarko as one of the biggest private companies in the world.

Figure II: Number of wells drilled by company

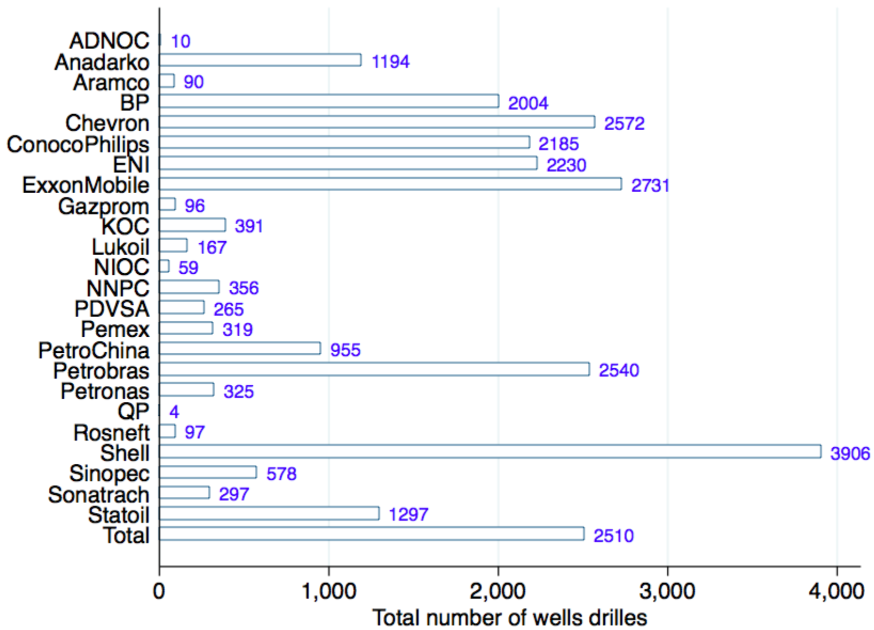

values with the group specific mean when calculating the company-location-quarter specific means.[5] Moreover, it is important to mention that we are completely missing information on US onshore drilling. However, our interviews with industry specialists suggest that US onshore drilling activity (in terms of activity and costs) is an outlier and, thus, should have been accounted for if we are interested in global drilling activity rather than activity in the US only.

We use the data to construct two variables. *Drilling activity*, which is the total

[5]The most important predictor of variable cost is the depth of the well drilled. In our sample the correlation coefficient is 0.46 explaining nearly half of the variation. Being aware of the positive correlation we can get an idea about the distribution of the missing variables by using the available information. In Figure VI (see Appendix) we plot the distribution of the depth of the wells for observations where costs are missing and where costs are not missing. Only in 125 cases the information is not available for both variables, which represents less than 0.5% of our sample. Comparing medians and means of the corresponding distributions we cannot reject the hypothesis of no significant differences. The similarity of distributions and the insignificant difference of means and medians suggests that, for our purposes, we might treat the missing observations as missing at random.

number of exploration and appraisal wells (henceforth exploration wells) drilled by a particular company in a particular location (onshore/offshore) in a particular quarter[6]. *Costs of drilling*, which represents the variable costs[7] of drilling in $US. Unit costs of drilling are transformed into real values using US CPI[8].

For the subsequent empirical analysis we transform the data as follows. We want to take into account location and company fixed effects in both dimensions: activity and costs.[9] Moreover, we need to take into account that these variables are not stationary. In order to do this we partly follow Kilian (2009) in his approach. We construct quarterly aggregates to capture drilling activity and costs of drilling on location[10] and company level (sum for drilling activity and mean for costs). The constructed time series is then logged and first differenced to account for the individual fixed effects. Finally, the individual series are weighted according to the number of wells drilled relative to the total number of wells drilled in the particular quarter. The resulting time series are presented in the second and third panels of Figure III.

In the first panel of Figure III the percentage change of the real oil price is presented. The nominal monthly oil price for WTI (Cushing) is taken from EIA. The oil price

[6]We have information on the starting day of drilling and on the completion day of drilling. We use the former because we are interested in how quickly companies in the oil and gas sector respond to an increase in the crude oil price.

[7]We note that the calculation of costs per well can vary depending on the region and the company involved. However, there is a great deal of effort from company's employees to construct comparable variable costs of drilling and exclude "back office" costs.

[8]Quarterly CPI for the US is taken from OECD statistics.

[9]We initially considered accounting for country fixed effects. For two reasons we have decided against this approach. First, the increasing sparsity of data points when disaggregating further. Second, we think of the companies as being globally integrated. Thus, variable cost (not taking into account the required negotiations over licences and the instability of countries) is driven by the technological abilities of firms, the location drilled, and the depth of the hole drilled rather than national borders.

[10]Onshore/Offshore

Table I: Descriptive statistics of main variables (weighted growth in percent)

Variable	Mean	Median	SD	Min	Max
Real oil price	1	3	14	-68	31
Number of wells drilled	12	12	13	-20	40
Costs of drilling	1	0	14	-20	47

Figure III: Time series of the main variables (weighted growth in percent)

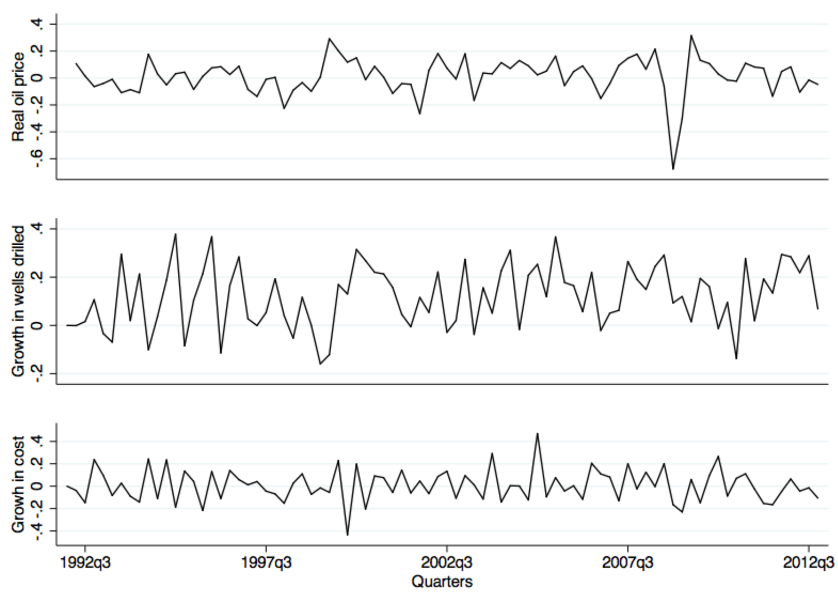

is transformed into quarterly values using the arithmetic mean, and into real values using US quarterly CPI.[11] Descriptive statistics of our variables are presented in Table I.

[11] Quarterly CPI for the US is taken from OECD statistics.

4 Identification and Estimation

Based on the theoretical section, we set up a structural model using quarterly data consisting of $Y_t = (\Delta ln(p_o), \Delta ln(\boldsymbol{D}), \Delta ln(\boldsymbol{c_D}))'$: the growth in the real price of oil, $\Delta ln(p_o)$; the (weighted) growth in the number of wells drilled, $\Delta ln(\boldsymbol{D})$; the (weighted) growth in costs of drilling a well, $\Delta ln(\boldsymbol{c_D})$.[12] The structural VAR representation is

$$\mathbf{A_0 Y_t} = \alpha + \sum_{\mathbf{i=1}}^{\mathbf{7}} \mathbf{A_i Y_{t-i}} + \varepsilon_{\mathbf{t}} \tag{1}$$

α is a vector of constants capturing the average growth rate. $\mathbf{A_i}$ is a matrix of the respective coefficients in period $t-i$. $\varepsilon_{\mathbf{t}}$ is a three-dimensional vector with serially uncorrelated and mutually uncorrelated errors. We assume that matrix $\mathbf{A_0}$ has a recursive structure such that the reduced form errors $\mathbf{e_t}$ can be decomposed according to $\mathbf{e_t} = \mathbf{A_0}^{\mathbf{-1}}\varepsilon_{\mathbf{t}}$:

$$\begin{pmatrix} e_t^{p_o} \\ e_t^{D} \\ e_t^{c_D} \end{pmatrix} = \begin{pmatrix} a_{11} & a_{12}=0 & a_{13}=0 \\ a_{21} & a_{22} & a_{23}=0 \\ a_{31} & a_{32} & a_{33} \end{pmatrix} \begin{pmatrix} \varepsilon_t^{p_o} \\ \varepsilon_t^{P_s} \\ \varepsilon_t^{P_d} \end{pmatrix}$$

The first structural shock, $\varepsilon_t^{p_o}$, is an *oil price demand shock* which is defined by unpredictable innovations to the oil price. We assume that the oil price does not respond contemporaneously (within a quarter) to shocks originating in the upstream sector, shocks to drilling activity or cost of drilling. This assumption is justified by the fact that increased activity in drilling exploration wells does not immediately translates into a change in oil and gas supply. It takes much longer than a quarter, and typically more than a year, before production starts after exploring and assessing new reserves. Alternatively, one could argue that expectations of market

[12]For a discussion on the identification of shocks and estimation of VARs see Kilian (2011)

participants about new discoveries effect the oil price. But it usually takes more than a quarter before a reasonable estimate of the size of an oil field is formed, and it takes even longer before total drilling costs can be determined (Adelman, 1990). Formally, we assume a_{12} and a_{13} to be zero.

The remaining variation in the reduced form errors - after controlling for the variation in the oil price - is decomposed into $\varepsilon_t^{P_s}$, a *supply shock*, and $\varepsilon_t^{P_d}$, a *drilling specific demand shock*. These are productivity shocks to firms which are typically thought of as being driven by technological and regulatory changes. Technological changes may be driven by advances in deep sea offshore drilling, improvements in seismic technology or development of horizontal drilling. Regulatory changes may be thought of as being an increase in taxes, tougher safety regulations, the expansion of area available for leases or the nationalization of an oil and gas industry.

Obviously, firms on the supply side (e.g. Halliburton, Schlumberger, Baker Hughes, etc.) and on the demand side (e.g. BP, Shell, Chevron, etc.) may be affected by productivity shocks. Thus, we refer to the shocks as a *supply shock* if it effects firms on the supply side and as a *drilling specific demand shock* if it effects firms on the demand side of the upstream sector. The latter is referred to as a 'drilling specific' demand shock to emphasize that it is a demand shock which originates in the upstream sector and is not driven by fluctuations in the oil price. To disentangle the two we assume that the supply curve for exploration drillings is vertical in the short run (within a quarter). Hence, only supply shocks can affect the number of wells drilled and costs of drilling contemporaneously. More formally, we assume a_{23} to be zero. This assumption requires a discussion.

Overall supply of engineers and rigs on the global level is limited in the short run: most certainly it requires more than a quarter to educate an additional engineer (Financial Times, 2014); from industry specialists we know that it typically requires more than a quarter to build a small rig and up to 3 years to build an offshore rig. One could still argue that supply can adjust at the intensive or the extensive margin in the presence of spare capacities. Kellogg (2011) reports that the scope for adjustments at the intensive margin is rather limited as hired drilling crews typically operate 24 hours a day and 7 days a week, rotating crews in 8 hours shifts. It is more difficult to argue that there are no spare capacities which would allow for an adjustment at the extensive margin. Ishii (2011) studies the oil and gas industry in the US and reports the existence of spare capacities which are, however, costly to activate. This suggests the existence of a steep but not vertical supply curve of drilling services. However, he focuses on all types of drilling rigs, including rigs used to drill production wells, in well explored oil and gas fields in the US. We, on the other hand, focus on the subsample of exploration wells for all the countries outside the US. Thus, we argue that oil companies exploring new territories outside well developed oil and gas fields face vertical supply curves in the short run. Mainly because rigs have to be utilized, or even constructed, transported to a new region, and a new crew has to be hired and possibly educated to operate the rig in the new territory. Our interviews with energy specialists suggest that this process requires more than a quarter, and usually up to 3 quarters.

Before estimating the VAR model we conducted the usual specification tests on the stability of the VAR, the normality and the serial correlation of the errors and the optimal number of lags. In order to achieve normality and no serial correlation of the errors we introduced a dummy variable to capture the negative oil price shock in

Figure IV: Responses to One-Standard-Deviation Structural Shocks

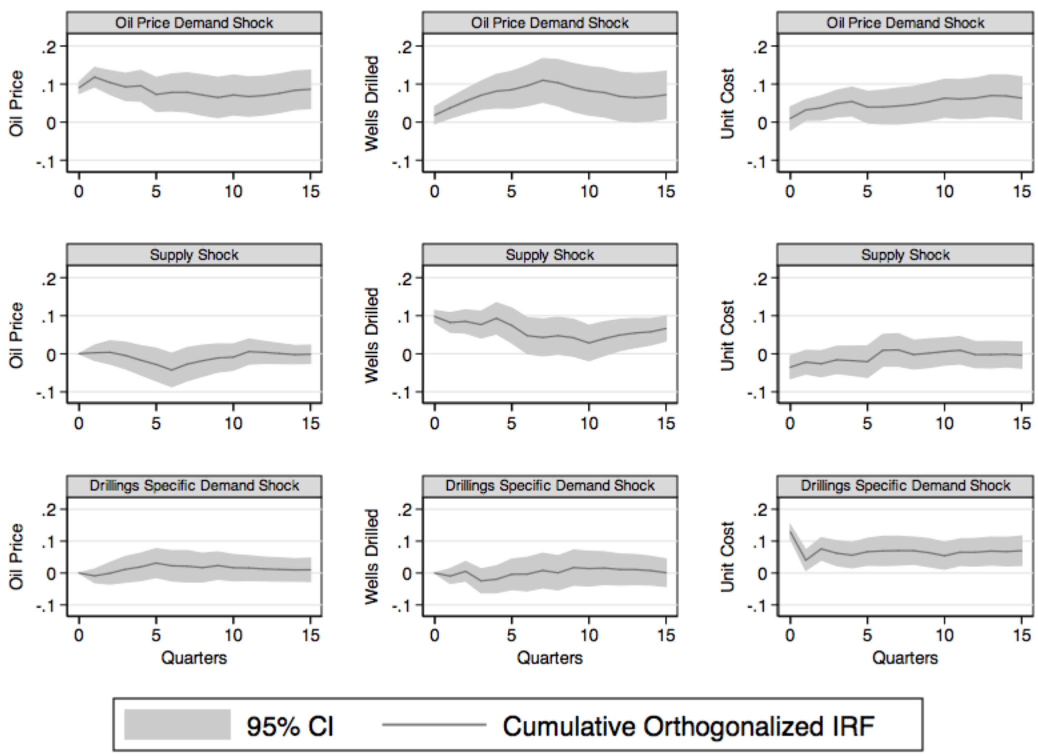

the last quarter of 2008 (see first panel of Figure III). Starting with 8 lags (2 years) the LR test suggests the use of 7 lags which we employed in our preferred specification. Finally, instead of relying on asymptotic theory we bootstrap the standard errors.

5 Results

The impulse response functions of our baseline results are presented in Figure IV. Following a structural *oil price demand shock*, $\varepsilon_t^{p_o}$, the number of wells drilled and average cost of drilling significantly increase from the second quarter onwards. Following a 1% increase in the oil price, the number of wells drilled increases by 1% and average cost of drilling increases by 0.5% after 3 to 4 quarters. Note that the increase is permanent for both variables.

Following a structural *supply shock* , $\varepsilon_t^{P_s}$, costs of drilling decreases. In particular, a 1% increase in drilling activity decreases costs of drilling by 0.25%. However, the effect on costs of drilling remains predominantly insignificant. This suggests that the demand curve is relatively flat. There wis a negative effect on the price of oil after approximately 6 quarters. However, the effect remains transitory and turns insignificant after approximately 2 years.

Following a structural *drilling specific demand shock*, $\varepsilon_t^{P_d}$, the effect on the number of wells drilled remains small and insignificant. The effect on the real price of oil is hump-shaped but insignificant on the 5% significance level. As in case of the structural *supply shock* the effect on the real price of oil does not appear to be permanent.

Note that the responses of drilling activity and costs of drilling to oil price demand shocks and drilling specific demand shocks differ. This result merits a discussion. Our interviews with industry specialists suggest that oil companies operate under relatively tight budget constraints. Essentially, this is driven by the well known link between debt ratios (e.g. total debt/total assets) and interest rates paid on capital employed in the production process (see for example Merton (1974)). Because the activity of oil companies is traditionally capital intensive and exploration is subject to great uncertainty this link is particularly important in the oil and gas industry (Adelman, 1962). A fact investors are very well aware of (Dumont, 2014).

The existence of a tight budget constraint suggests that drilling activity and cost of drilling will react differently to different demand shocks. A relaxation of the budget constraint following an oil price driven demand shock suggests that oil and gas companies get involved in already prescheduled projects which have not been approved in the past. On the other hand, an increased demand for drilling services following a drilling specific demand shock due to technological innovations and regulatory changes will be usually matched by a decrease in drilling activities in other regions due to tight budget constraints. Following a similar line of reasoning, projects initiated by oil companies following an oil price demand shock are more likely to take place in already well explored territories such that changes in drilling costs may remain within limit. On the other hand, drilling specific demand shocks will be typically the result of technical advances or regulatory changes allowing the exploration of a new region which are most certainly subject to a vertical supply curve. Summing up, we argue that drilling specific demand shocks and oil price demand shocks result in a different combination of projects which in turn have a different effect on drilling activity and costs of drilling.

6 Conclusion

We use micro data from Wood MacKenzie to construct two quarterly time series capturing the total number of exploration wells drilled by the 25 major oil and gas companies and (ii) the average cost of drilling. In combination with the real price of oil we estimate a three dimensional structural VAR model. The proposed model allows us to decompose the variation in the reduced form errors from estimating the VAR into three structural shocks by assuming a recursive structure: an *oil price* *demand shock*, a *supply shock* and a *drilling specific demand shock*. We estimate the dynamic effects of these shocks on drilling activity, costs of drilling and the real price of oil.

We have three main results. First, our results suggest an upwards sloping supply curve in the upstream sector of the oil and gas industry with a long run price elasticity of supply of 2. Second, supply shocks and drilling specific demand shocks do *not* permanently affect the oil price. This is particularly interesting because it calls into question the theoretically expected link between extraction costs and the price of oil (Krautkraemer, 1998). We follow Kilian (2009) in arguing that this result may be explained by the endogenous expansion of crude oil production as a reaction to an increase in drilling costs.

References

Adelman, M. A. (1962): *The Economics of Petroleum Supply: Papers by MA Adelman, 1962-1993*. MIT press.

Adelman, M. A. (1990): "Mineral depletion, with special reference to petroleum," *The Review of Economics and Statistics*, 72(1), 1–10.

Anderson, S. T., R. Kellogg, and S. W. Salant (2014): "Hotelling Under Pressure," Working Paper.

Dumont, C. (2014): "4 Leverage Ratios Used In Evaluating Energy Firms," `http://www.investopedia.com/articles/fundamental-analysis/12/4-leverage-ratios-used-in-evaluating_energy_firms.asp`, Accessed: 2014-08-03.

Financial Times (2014): "Terrifying oil skills shortage delays projects and raises risks," July 16.

Hamilton, J. D. (2008): "Understanding crude oil prices," Discussion paper, National Bureau of Economic Research.

Helman, C. (2013): "The World's 25 Biggest Oil Companies," Forbes.

Ishii, J. (2011): "Useful Excess Capacity? An Empirical Study of US Oil & Gas Drilling," Working Paper.

Kellogg, R. (2011): "Learning by drilling: Interfirm learning and relationship persistence in the Texas oilpatch," *The Quarterly Journal of Economics*, 126(4), 1961–2004.

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

——— (2011): *Structural vector autoregressions*. CEPR.

Krautkraemer, J. A. (1998): "Nonrenewable resource scarcity," *Journal of Economic Literature*, pp. 2065–2107.

Merton, R. C. (1974): "On the pricing of corporate debt: The risk structure of interest rates*," *The Journal of Finance*, 29(2), 449–470.

Venables, A. J. (2011): "Depletion and development: Natural resource supply with endogenous field opening," Discussion paper, CESifo working paper: Resource and Environment Economics.

Yergin, D. (2011): *The prize: The epic quest for oil, money & power*. Simon and Schuster.

7 Appendix

Figure V: Number of wells drilled

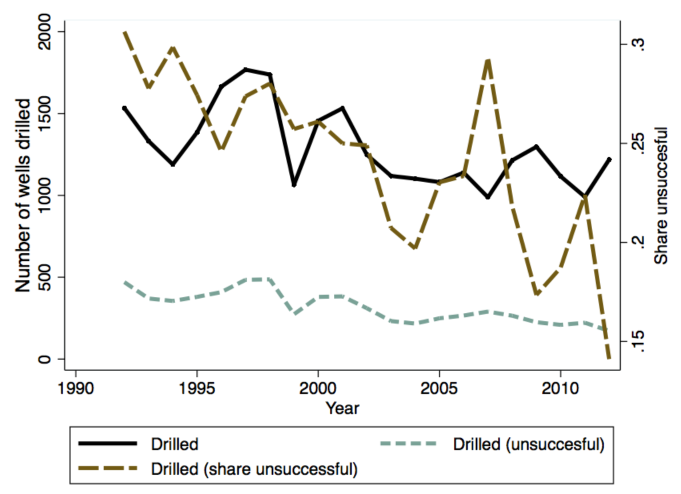

Figure VI: PDF of total depth (metres, logged)