

The relationship between oil prices and rig counts: The importance of lags

Ahmed Khalifa¹

Corresponding author
College of Business and Economics
Qatar University
2713, Doha, Qatar
aliabdelkh@qu.edu.qa

Massimiliano Caporin

Department of Economics and Management “Marco Fanno”
University of Padova
Via del Santo, 33
35123 Padova, Italy
massimiliano.caporin@unipd.it

Shawkat Hammoudeh

Lebow College of Business
Drexel University
Philadelphia, PA, United States
IPAG Lab
IPAG Business School
France
hammousm@drexel.edu

¹ Corresponding author, College of Business and Economics, Qatar university, Doha, Qatar, Email: aliabdelkh@qu.edu.qa, Tel: +974-4403-6498, Fax: +974-4403-5081

Abstract

This study deals with a timely and relevant issue in the oil market, which is the relationship between changes in oil prices and changes in rig counts in the wake of the drop in oil prices, while accounting for other determinants of this relationship. This relationship is of strong interest to analysts, investors and policymakers in the United States and other countries. We empirically verify the impact of changes in oil prices on the rig counts, which has a lag of up to one quarter. This evidence is stable across time and over different linear and non-linear models. The analysis also suggests that the relationship is non-linear, which is verified by both the quantile regression and quantile-on-quantile models. We find evidence of non-linearity that has softened in the most recent years where the relationship between the variables has stabilized.

JEL: C22, C58, G11, G15, Q31, Q41.

Keywords: Rig counts, Oil price, Lags, Quantile regression, Quantile-on-Quantile

Executive summary

The study deals with a timely and relevant issue in the oil market, which is the relationship between changes in oil prices and changes in rig counts. This relationship is of significant interest for analysts, investors and policymakers, whether they are oil companies, commercial banks, investment banks, in the United States and other countries. We verify empirically the existence of a positive impact from lagged oil returns (up to one quarter) to changes in rig counts, while taking into account the influence of other pertinent factors. Moreover, we show that this relationship changes over time and is likely to be more complex than what can be captured by a simple linear model. By observing the results of several models, we find a common denominator where all the methods used show that oil returns positively affect changes in rig counts, as is documented by the practitioners' literature, with a lag of up to three months. This implies that changes in oil prices are not immediately reflected in changes in the rig counts. Further, we note that the relevance of the price-rig relationship changes over time but becomes stronger and more stable from 2005 onward, even with controlling for the potential impact coming from other economic and financial variables. Further, the non-linear models show that the impact of oil returns on changes in rig counts becomes stronger when the oil returns take large negative values. Therefore, the oil impact is stronger during bearish oil markets. This result may have an implication for the recent 70% plunge in the oil prices. It implies that the ensuing fall in oil rig counts and the impact on oil prices will be large, but the result will come with a lag. In turn, this has strong consequences for the prevailing oil glut over time. Since the literature shows that changes in oil rig count affect production, then this result is relevant for the duration of the oil glut as it must be shorter than the market expects, or it may point to a significant drop in U.S. oil production but with a lag

1. Introduction

Oil is the most global volatile commodity with repercussions reverberating throughout the global oil industry and the world economy. Oil prices spiked from a low of about \$3.6/barrels in 1972 to a peak of more than \$145/barrels in 2008. They then collapsed to reach a nadir in early 2009 due to the recent global financial crisis but then recovered to exceed \$100/barrels by the middle of 2014. More recently, they again plunged by about 70% due to the boom in shale oil and the geopolitical changes toward Iran. The fluctuations in oil prices have strong impacts on rig counts and drilling activity (investments) and well productivity. For example, the recent fall in oil prices accompanied with a drop in the drilling activities, which plunged from the peak of 1840 rigs in Dec. 29th to 541 rigs on February 12th 2016 (a 71% fall)¹.

A rigs is a machine that rotates the drill pipe from surface in order to drill a new well (or sidetracking an existing one) to explore for, develop and produce oil.² Although, economic theory and a large body of the existing literature on this topic highlight that a higher price will stimulate investment and vice versa (Kellogg, 2014), the relationship between changes in oil prices and rig counts may be not that obvious and may not also be direct because of the presence of lags (Black and LaFrance, 1998). It is possible for rig counts to continue to change, while oil production increases and oil prices drop.³ Rig counts can also go silent, while production is going on as happened in North Dakota recently. The relationship thus may not be linear because

¹ Source; <http://www.aogr.com/web-exclusives/us-rig-count/2014>

² In general, wells can be one of three types: exploratory, development, or infill. For the details about the three types of wells, see (Kellog ,R. 2014).

³ During the great oil crash of 2008-2009, oil peaked at \$145 a barrel in the week that ended on July 11, 2008. On the other hand, the rig count kept moving up until the week that ended on November 7, 2008, or 119 days later, despite the fact that oil had plunged 50% from its highs by then. Once the price of oil bottomed out at \$34/barrel in December of 2008, the oil rig count continued dropping until the week that ended on May 22, 2009, 154 days later. Recently, the price of oil peaked in June 2014, while the rig count peaked 112 days or approximately four months later, that is during the week of October 10, 2014 approached 1604 rigs. Similarly, the price of oil bottomed around March 13, 2015, while the rig count bottomed round June 26, 2015, or 105 days later, it approached 629 rigs. See Majeure (2015).

of the lags as well as the effects of changes in oil well productivity, rig efficiency, drilling costs, commodity inflation, hedging, changes in inventories, etc. (Hunt and Ninomiya, 2003). In this case, quantifying this relationship would require the use of non-linear models.

There are several reasons that explain the presence of lags in rig drilling. During periods of lower oil prices, oil companies initially revisit their resources that are reckoned to be uneconomic. There are also rig contracts and rigs rented for a number of years, which stand in the way of suddenly terminating drilling activity. The lags are also present during higher oil periods as it takes time to acquire new leases/concessions, carry out seismic surveys, recruit workers, etc. (Ghoury and Aneesuddun, 2015). The presence of lags also reflects the oil exploration and production (E&P) companies' efforts to make sure that there is a forming trend and a demonstrating longevity in the oil market that warrants increasing or decreasing the drilling activity. Adding new drills (investment decisions) may also require financing from banks, which may not be forthcoming right away because those financial institutions want to make sure that the oil reserves collateral is worth the risk. There is also a gestation period that characterizes capital investments as it takes time to move the rig to its new drilling location and to transport work force and equipment to the drilling site (Osmundsen et al., 2008). On the down side, if oil prices start to fall, the rigs in progress will continue in drilling until the wells are completed because it is expensive to shut down a rig once it started due to installation costs and the rent contract.

The objectives of this paper are fourth fold. First, is to provide further evidence of the relationship between oil returns and changes in rig counts, while accounting for relevant economic and financial variables. Second, which is the initiator or the leader of this relationship? Third, if the relationship is not contemporaneous, then what are the exact timing, the potential

non-linearity, the time variation and the structural changes in this relationship? Fourth, in which market or economic conditions is this relationship stronger or weaker, given the fact that oil prices go through boom, normal and bust periods? Realizing these objectives is important for the policymakers and analysts. If it is the former, then we know that bullishness in the oil market will have an impact on the oil E & P companies, which seek loans from banks by offering their oil reserves as collateral. If it is the latter, it may be a curser that changes in drilling activity could presage higher oil prices.

It is important to discern whether the initiator of the relationship between changes in oil prices and rig counts is oil demand or oil production as a starting point in the empirical investigation of the paper. The change in oil demand is usually incremental and takes time, which means longer lags and a smooth and steady change in rigs. If the cause of the lag is oil production, then the lags may be different, the directional relationship may be inverted and the change in rigs may be abrupt.

The investigation of the relationship between changes in oil prices and rig counts will be explored over the period January 1990-July 2015. It will use both the quantile regression analysis (QRA), as developed by Koenker and Basset (1978), and Koenker (2005), and the quantile-on-quantile (QQ) approach, advanced by Sim and Zhou (2015), to measure the relationship under normal and extreme conditions, which fits the erratic behavior of oil prices. We have chosen these techniques for different reasons. First, because they can address changes in relationships over bullish, normal and bearish periods of the dependent variable. In fact, those techniques allow for separating the dependence between variables in the upper tail from the dependence in the median or the lower tail. In contrast, the linear models analyze only mean dependence and not that of the tails. Consequently, the quantile methods are more flexible than the standard

linear regression model, and represent an efficient approach for detecting the presence of interdependence asymmetries between the analyzed data across quantiles. Further, those methods allow for relationships across the variables that are specific to the location of the dependent variable's observations over its density support (QR), or are specific of the location of the dependent variable and of one of its relevant covariate over their respective density supports (QQ). This allows for making a cascade of analyses starting from the linear models, moving to QR and then to QQ with an increasing model complexity, and also allowing for more relationships affected by the location, the scale and the shape of either the response variable and/or the covariates. Finally, QR allows identifying shifts in the propagation mechanism across variables when shocks hit the dependent variable, whether QQ extend this to the occurrence of shocks on both the dependent variable and one relevant covariate.

Our results show that a directional relationship exists between changes in oil prices and rig counts but with some lags. This is consistent with the practitioners' literature which estimates the lag to be between 2 and 3 months, and coherent with other preliminary analyses which highlight symptoms of such a lagged relationship. Furthermore, we show evidences supporting the instability of the relation over time, and across quantiles of the change in rig counts. This suggests that both market and economic conditions might influence the lagged relationship between rig counts and oil price movements. Notably, we observe that changes in rig counts react more to oil returns when the change in the rig counts takes values below its median or is not very strong, which is consistent with the previous literature such as Kellogg, 2014 who finds that the response of drilling activity to changes in price volatility is commensurate with the predictions of the real options theory. Furthermore, when analyzing at a finer detail or a nuance of the dual relationship, that is, by conditioning both on the quantiles of the changes in the rig

count and oil returns, we show that the impact of oil returns on changes in rig count is much higher when the oil returns take on very negative values. Therefore, downside movements in oil prices lead to larger decreases in rig counts. Furthermore, by contrasting this result over two different sub-samples, we note that the large impact of the large negative oil returns on rig counts decreases in the most recent years.

Our results are important to investors in the oil markets who will benefit from this knowledge of lags as they decide on whether to long or short the shares of the oil companies before they spike or plunge in reaction to changes in rig counts. However, this useful information will not be complete until analysts and policymakers acquire information on the (speed of) adjustment of supply to price changes. The results we provide will be also helpful to foreign direct investments flowing to the oil industry. Methodologically, the quantile framework that the paper seeks to build should help improve forecasting the impact of oil price returns on the rig count.

The remainder of this study is organized as follows. Section 2 discusses the related literature. Section 3 describes the data and provides preliminary analyses. Section 4 introduces the research methodology. Section 5 presents the empirical results, and Section 6 concludes the study.

2. Related literature

There are two strands of the literature that deal with drilling activities. The first strand focuses on the determinants of the investment in the drilling activities and the second examines the relationship between changes in oil prices and rig activity, with an attention to lags. Within

the first strand, Anderson et al. (2014) provide evidence that the production of oil from existing wells in Texas does not respond to oil price incentives but drilling activity and costs respond strongly to oil prices. The authors reformulated the Hotelling theory of exhaustible resources in the form of a drilling problem to where firms choose when to drill, but the decaying production from existing wells is constrained by reservoir pressure, to arrive at their results.

Using a dynamic model of firms' investment problem, Kellogg (2014) estimates the effect of lagged price changes on drilling activity (rig counts,) suggests that the major impact of price changes on drilling activity (rig counts) occurs after 3 months and shows that a study on the determinants of drilling activity (rig counts) should take volatility as a determining factor into account as well. Due to shortages in drilling contracts and personnel, Osmundsen et al. (2008) discusses designing incentive contracts in the drilling sector. This issue is represented by the compensation formats utilized in the present and in the consecutive drilling contracts. The author finds that changes in contract format pose a number of relevant questions relating to resource management. The questions include the following. Do evaluation criteria for awarding drilling assignments boost the development of new technology and solutions? Finally, how will a stronger emphasis on drilling efficiency influence reservoir operation?

With respect to the oil price deterioration and its consequences on the investments returns in the energy drilling, Toews and Naumov (2015) estimate a VAR for the oil and gas upstream industry and annual data to examine the dynamic effects of oil price and drilling activities. These authors find a directional relationship that runs from changes in real oil prices to rig counts with a one-year lag. They show that a 10 % increase (decrease) in the real oil price causes a 4% increase (decrease) in the global drilling activity and a 2% rise in the cost of drilling with a lag of 4 and 6 quarters, respectively. They also find that positive shocks to drilling activity

affect the oil price negatively; however, the shocks to costs of drilling do not have a permanent effect on the price of oil. Within the same strand of literature, Henriques and Sadorsky (2011) investigate how oil price volatility affects the strategic investment decisions of a large panel of U.S. firms. They use key insights from the real options literature to develop a model of a company's strategic investment and show how changes in oil price volatility can affect strategic investment decisions. Their model is estimated using the generalized method of moment estimation techniques for panel data sets. They show that there is a U-shaped relationship between oil price volatility and firm investment. Their results are consistent with the predictions from the strategic growth options literature.

The second strand of the literature relates changes in oil rig activity to changes in oil prices, and underscores the importance of lags. Ringlund et al. (2008) estimate relationships between oil rig activity and crude oil prices in different non OPEC countries, using dynamic regression models. These models are augmented with latent components capturing trend and seasonality. The authors show a positive relationship between rig activity and oil prices, but the strength of the relationship differs across the different regions, depending on the oil industry structure and the reaction of oil rig activity to oil price changes. Overall, their results show a clear relationship between the oil industry structure in the region and the reaction of oil rig activity to price changes. The authors also find that the long-run price elasticity for oil rig activity in non-OPEC countries on average is around unity.

Osmundsen and Mohn (2006) and Mohn and Osmundsen (2008) find muted short-term effects but robust long-term effects of oil prices on exploration activity in the Norwegian Continental Shelf during the period 1965-2004. OGJ (2003a) demonstrates figuratively that changes in oil rig activity in the United States tend to follow but with lag the fluctuations in oil

prices. On the other hand, Abraham (2000) contends that the oil industry has been quick to boost drilling activity whenever oil prices have remained high for a minimum of six months, at least until the year 2000. Black and LaFrance (1998) question the lag based on an empirical investigation of oil fields in Montana.

Some literature relates oil rig activity to economic activity. Brown (2014) among others estimates the response of total employment in oil- and gas-producing states to changes in rig activity caused by changes in oil prices. The author finds that eliminating an active rig results in job losses in the long run.

To our knowledge, the current study will be the first to examine the relationship between changes in oil price returns and changes in rig count using both the QRA the Q-Q approach. It pays considerable attention to lags, which have been acknowledged in previous literature but has not studied with the same vigor and rigor we express in this work.

3. Data Description

The two most relevant variables in our study are the oil price (WTI), and the rig counts monitored at the US level⁴. We also include in our analyses a set of potentially relevant economic and financial covariates that could affect the evolution of either the oil price, the rig counts⁵ or both. Specifically, we consider the following additional variables. First, we track the lagged rig productivity, ⁶, which reflects the evolution of both the oil prices and the rig counts.

⁴ All of the data series are obtained from Bloomberg and the Federal Reserve Bank of Saint Louis.

⁵ The Baker Hughes Rotary Rig count includes only those rigs that are significant consumers of oilfield services and supplies and does not include cable tool rigs, very small truck mounted rigs or rigs that can operate without a permit. Non-rotary rigs may be included in the count based on how they are employed. For example, coiled tubing and work-over rigs employed in drilling new wells are included in the count. To be counted as active, a rig must be on location and be drilling or 'turning to the right'. A rig is considered active from the moment the well is 'spudded' until it reaches target depth or 'TD'. Rigs that are in transit from one location to another, rigging up or being used in non-drilling activities such as work-overs, completions or production testing, are NOT counted as active. Miscellaneous rig counts represent geothermal rigs.

⁶ We have selected the US rig counts, because the United States is one of the largest oil producers in the world, is the global leader in oil rigs, bases investment decisions on rational criteria and has a good database.

We then consider two oil and economic variables, the oil inventory levels and industrial production which is associated with the business cycle and affects the demand for oil. As proxies for the world industrial production and global oil inventory level, we take the U.S. industrial production index (seasonally adjusted from Bloomberg) and the U.S. oil inventory level. Further, we consider the trade weighted dollar exchange index to monitor the effect of the reference currency that is used in oil pricing relative to global currencies. We then move to indicators of financial stress, and consider the financial condition index, the financial stress index, the VIX index and the MOVE index, which is the Merrill Lynch Option Volatility Estimate.

Finally, we consider variables monitoring the bond/credit market. We first include the effective federal funds rate and the 10-year US Treasury notes redemption yield. Then, we consider the TED spread (difference between the interbank loan rate and the Treasury bill rate) which is an acronym formed from T-Bill and ED, the ticker symbol for the euro-dollar futures contract, and the term spread between the 10-year Treasury note rate and short term rates as a predictor of real economic activity. The data for those variables were sourced from the database of Federal Reserve Bank of St. Louis, with the exception of the data for MOVE which was obtained from Bloomberg. The monthly sample we consider covers the period from September 1990 to June 2015, giving us a sample size of 298 months.

We provide here a limited descriptive analysis of the major variables that are most relevant for this study, namely, the oil price, the rig counts and the rig productivity. The plots and similar analyses for the other variables, the covariates that might affect oil price and rig count movements, are available upon request. For the three variables, we perform standard diagnostic tests for stationarity, which detect the presence of unit roots, and for cointegration

which suggest the absence of cointegration. Thus, the results of those tests point us to a framework that models the oil returns, the changes in rig counts and rig productivity together.⁷ Figure 1 plots the levels and the changes of these three variables which highlight the big movements in oil prices in 2008 and in 2015, the increase in the rig counts after 2010 and the drop in 2015, and the decrease in the rig productivity, which is particularly evident in the first and second quarters of 2015.

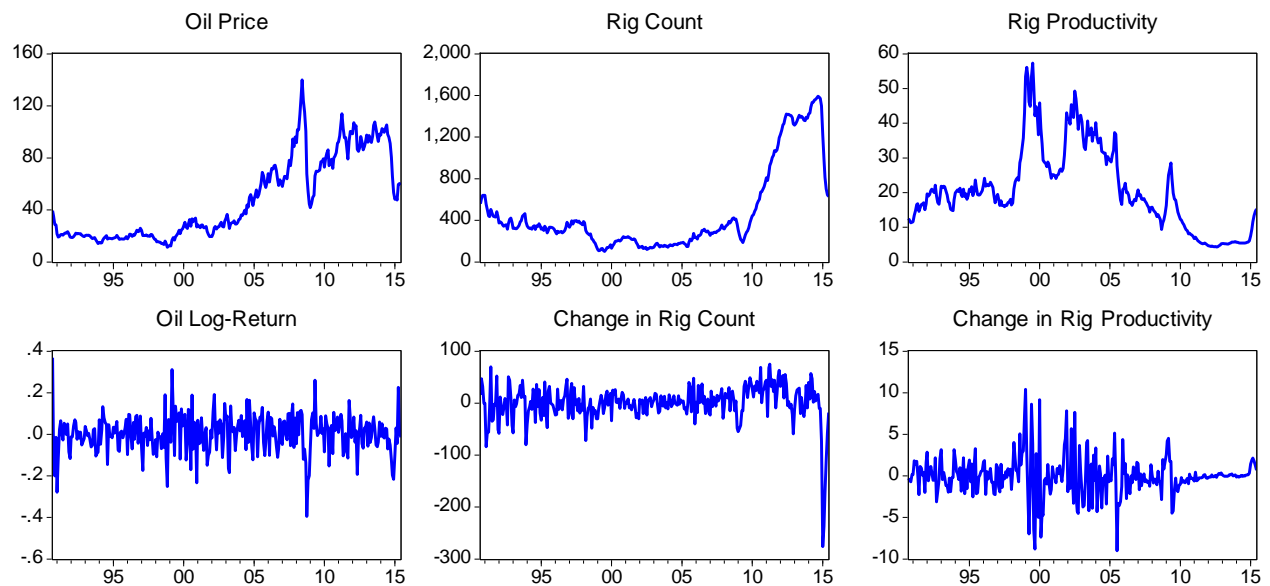


Figure 1: plots of Oil price, Rig count, Rig productivity, Oil log-returns, Change in rig count and Change in rig productivity.

For the other variables we consider in the study, we generally work on the variable changes (the first difference). The only exception is the industrial production index for which we consider the log-return.

To evaluate the presence of links across the three most relevant variables, we start by analyzing their contemporaneous correlation, see Table 1.

⁷ The tests' results and further descriptive analyses are available upon request.

Table 1; The contemporaneous correlation between Change in Rig count and change in rig productivity

Table 1: full sample correlations	Oil return	Change in rig count
Change in rig count	0.047	
Change in rig productivity	-0.054	-0.495




Figure 2: Rolling correlations across oil returns, rig count changes and productivity changes.

Notes: The plots, left to right, represent the correlations between the oil returns and the rig count changes; the oil returns and the rig productivity changes; and the rig count changes and rig productivity changes.

The full sample analysis shows that the contemporaneous correlations for oil returns with either changes in the rig count or changes in rig productivity are very weak. On the contrary, the rig count and rig productivity changes are strongly negatively correlated. Those correlation values may, however, change because of movements in the oil price, rig count and rig productivity in response to shocks or swings, as we observe, for instance, in the last part of the sample (see the lower panels of Figure 1). This is thus challenging the stability of the correlations across the entire sample. The same narrative applies to the variances of those variables, which might be characterized by large local movements. Based on these results for the full sample, we thus evaluate the volatility and correlations across the three most relevant variables (oil returns, changes in rig count and changes in rig productivity) over a 60-month rolling window. Figure 2 provides the graphical evidence. Notably, the relationship between the

changes in the rig counts and the oil returns oscillates around zero, and therefore we find a confirmation of the full-sample's results. Similarly, oil returns are almost unrelated to changes in rig productivity. Finally, changes in rig productivity and rig counts are negatively correlated, suggesting that an increase in rig counts is associated with a decrease in productivity. This negative relationship reflects the E&P companies' efforts to increase oil production to pay off debt and acquire new loans from banks. Moreover, there is inefficiency in acceleration in the energy industry.

Table 2: Cross correlations between oil returns and changes in rig counts

Lags	Changes in Rig Count versus Lags of Oil Returns	Oil returns versus lags of change in rig count
0	0.0475	0.0475
1	0.2333*	-0.0954
2	0.3116*	-0.1037
3	0.2783*	-0.1319*
4	0.2054*	-0.0768
5	0.1589*	-0.0869
6	0.1228*	-0.0499
7	0.0558	-0.0126
8	0.0284	0.0028
9	0.0113	-0.0405
10	0.0403	0.0177
11	-0.0416	-0.0097
12	-0.0089	-0.0374

Notes: The star * identifies the 5% statistically significant coefficients.

Thus, the results highlight an absence of a contemporaneous correlation between oil returns and changes in the rig count. However, we are still not aware of their possible lead-lag relationship, which is important for *analysts, investors and policymakers*. In order to shed some light on this causality, we provide in Table 2 the cross-correlations between the two variables. We have now an interesting finding, which is the following: the lags for the oil returns seem to affect the changes in the rig count, while the opposite relationship is very weak. As indicated earlier, this result is consistent with the scant practitioner and academic literature that has dealt

with lags. It also makes sense because the E&P companies want to make sure that their total revenues are increasing before they embark on new drilling.

The previous results suggest focusing on the use of dynamic models, but still we have evidences and methods designed to detect the presence of linear relationships. To get some insight into the possible presence of a non-linear relationship between changes in rig counts and oil returns, we consider the exceedance correlations measure proposed by Longin and Solnik, (2001). The intuition behind using this measure is to focus on the correlation across variables conditioning on a specific section of their joint density support, ideally in the tails. However, given that our main interest is in the impact of oil returns on the change in rig counts, we modify the exceedance correlation of Longin and Solnik (2001) as in Caporin et al. (2014), by conditioning only on the change in rig count quantiles. Therefore, we compute the following conditional correlation measures

$$\rho_{\tau}^{-} = \text{Corr}(\Delta Rig_t, \Delta \ln(Oil_t) | F(\Delta Rig_t) < \tau), \quad \tau \leq 0.5 \quad (1)$$

$$\rho_{\tau}^{+} = \text{Corr}(\Delta Rig_t, \Delta \ln(Oil_t) | F(\Delta Rig_t) > \tau), \quad \tau \geq 0.5 \quad (2)$$

where $F(\Delta Rig_t)$ is the cumulative density function (CDF) of the change in rig counts and τ is a pre-specified quantile. The exceedance correlation ρ_{τ}^{-} measures the association between ΔRig_t and $\Delta \ln(Oil_t)$ when the change in rig counts is located in its lower τ empirical quantile, with τ taking values below 0.5 that is below the median. On the other hand, ρ_{τ}^{+} monitors the correlation between the two variables when the change in rig counts takes values above the median, for τ taking values above 0.5.

Figure 3 plots the exceedance correlations for the contemporaneous change in rig counts and the oil returns, as well as for the cases where we lag the oil returns and where we focus on

the last quarter oil returns. Therefore, we compute the exceedance correlations $\rho_{\tau}^{-} = Corr(\Delta Rig_t, X | F(\Delta Rig_t) < \tau)$, with $X = \{\Delta \ln(Oil_t), \Delta \ln(Oil_{t-1}), \Delta \ln(Oil_{t-2}), \Delta \ln(Oil_{t-3}), \sum_{j=1}^3 \Delta \ln(Oil_{t-j})\}$, corresponding to the contemporaneous, lag 1, lag 2, lag 3 and the quarter cases (see Figure 3). These analyses on lagged oil returns and the use of the last quarter return stem from the cross-correlations where lags of oil returns up to one quarter seem to be more related to changes in the rig counts as opposed to the contemporaneous case. The exceedance correlation analysis allows us to verify whether the impact of oil prices on rig counts changes depending on the change in the rig count, being associated with increases in the rig counts (upper quantiles) or decreases in rig counts (lower quantiles).

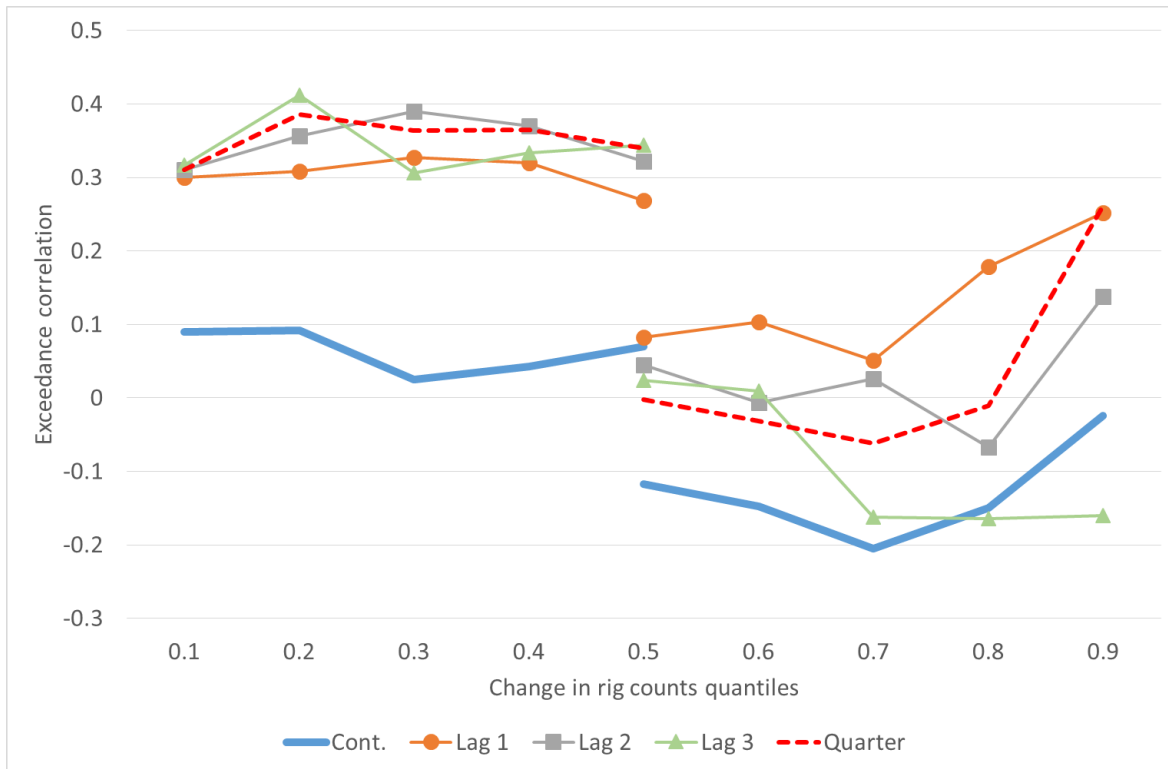


Figure 3: Exceedance correlation between the change in rig counts and oil returns for different quantiles τ of the change in rig counts.

Notes: On the left and up to the median (0.5 quantile) we report ρ_{τ}^{-} while from the median and above we report ρ_{τ}^{+} marking both cases with the same colors. In the blue solid line case, we have the exceedance correlation between the

changes in rig counts and the oil returns at time t , thus is the contemporaneous exceedance correlations. The exceedance correlations between the changes in rig counts and different lags of oil returns are represented by the orange circle line (1 lag), the gray square line (2 lags), the green triangle line (3 lags), and the red dashed line (last quarter).

The figure shows that the exceedance correlation significantly changes when moving from the case in which we analyze the association between contemporaneous variables to the cases where we lag the oil return to lags 1-3 and one quarter. In general, the use of lagged returns leads to an increase in the exceedance correlations, apart for the above median quantiles for the case of three periods' lags. Even more interestingly, there is a marked difference between the quantiles below the median and the quantiles above the median. In the first case, the relationship is positive and quite stable, and is not much affected by the lags between changes in rig counts and oil returns, the relevant element being the presence of a lag. In the other situation, where the quantiles are above the median, we have a more heterogeneous behavior across all the cases we consider, with exceedance correlations being close to zero for non-extreme quantiles, and taking positive or negative values when moving toward the upper tail. The marked difference between the exceedance correlations above and below the median suggests that the relationship between the variables is not simply linear. This implies that analysts, investors and policy makers should not expect to see a fixed relationship between changes in oil prices and rig counts, may be due to binding contracts, delays and differences in strength and rigor of the responses. Building on this evidence, we thus suggest the use of the quantile regression methods as a flexible tool to capture this potential non-linearity.

4. Research methodology

We report here a set of tools which we will consider to analyze the dynamic relationship between changes in oil price (returns) and changes in rig counts. In a first step, we allow for a simple dynamic linear structure. Let $Y_t = [\Delta \ln(Oil_t) \ \Delta Rig_t]'$ denote the two-component vector of our target variables and X_{t-1} be a k -dimensional vector of lagged control covariates. We specify the following Vector Auto Regressive model with “eXogenous” variables (VARX)

$$Y_t = \mu + \sum_{j=1}^p \Phi_j Y_{t-j} + \sum_{i=1}^q \beta_i X_{t-i} + \varepsilon_t. \quad (3)$$

Note that there exist several methods and approaches to identify the optimal orders p and q . However, to keep a balance between the model flexibility and the economic interpretation of the outcomes, we fix $p=3$ and $q=1$. The estimation of the model can be easily achieved by the least square methods, even with accounting for heteroskedasticity and autocorrelation in the residuals. A full-sample estimation of the VARX model in Eq. (3) allows one to determine whether there is a long-term dynamic relationship between the two target variables in the presence of the control covariates. Furthermore, the estimates open the door to the Granger causality testing. However, the dynamic relationship in Eq. (3) might not be stable over time. Therefore, we will also focus on a rolling estimation of Eq. (3) as well as on testing for the presence of structural breaks in this equation.

The dependence between changes in the rig count and oil returns may however be non-linear or not restricted to dependence across their (conditional) means. Therefore, in order to be able to detect a possible presence of causality in a more general framework, allowing for dependence across quantiles, we consider the estimation of a quantile regression model. We denote the τ conditional quantile of a target variable y_t (either oil returns or changes in the rig count) as follows

$$Q_\tau(y_t) = \delta_{0,\tau} + \delta'_{1,\tau}Z_t \quad (4)$$

where Z_t may include lags of both the target variables and lags of the control covariates. The coefficients included in δ_1 in Eq. (4) allow for detecting the presence of the impacts from the covariates (either target variables or control variables) at the quantile of order τ . To estimate the model, we minimize a criterion function based on the asymmetric loss

$$\rho_\tau(X) = X(\tau - I(X < 0)) \quad (5)$$

where $I(\cdot)$ is the indicator function. Koenker and Basset (1978) have shown that the minimization of the expected asymmetric loss in Eq. (5), with a proper linear function replacing X , allows for estimating the conditional quantiles. With respect to the model in Eq. (4), this way corresponds to the following minimization problem

$$\min_{\delta_{0,\tau}, \delta_{1,\tau}} \sum_{t=1}^T \rho_\tau(y_t - Q_\tau(y_t)). \quad (6)$$

Clearly, the model in Eq. (4) provides estimates of the coefficients that are dependent upon the chosen quantile τ , thus motivating the indexation of the coefficients on quantiles. We refer the readers interested in further details to Koenker (2005). Among the various inferential procedures available for the estimates of the quantile regression, we will consider first the evaluation of the significance of the coefficients. This follows the standard tools as the estimates have an asymptotic Gaussian density (Koenker, 2005). Further, and more relevant for the purposes of this paper, we will test the stability of the coefficients across different quantiles. In fact, the coefficients that are equal across a number of selected quantiles support the so-called pure location-shift hypothesis, where covariates do have the same impact across all quantiles, and therefore a linear model would be appropriate (instead of a more flexible quantile regression model). The quantile stability test can thus allow for answering a first question relating to the

presence of non-linear relationships between the target variables and the control covariates. If a non-linear behavior is present, the estimated coefficients can be analyzed to detect any differences in the impacts of the covariates across quantiles.

As in the linear VARX specification, we will estimate the quantile regression model both on the full sample as well as by focusing on the rolling methods. The rolling analyses will allow one to verify the stability of the coefficients over time. The quantile regression methods evaluate the dependence of the quantile of the target variable on a set of control covariates. However, the relationship may change depending on the location of the covariates with respect to their own density support. Two approaches are viable which are: the estimation of a multivariate multi-quantile system as in White et al. (2015), where quantiles of different variables possibly interact, and the simpler and more manageable quantile-on-quantile approach of Sim and Zhou (2015). We opt for the second approach for two main reasons. First, we are interested in the dependence among the quantiles of two variables, while also accounting for the presence of control covariates (not included in the work of White et al., 2015). Second, the presence of a somewhat limited sample size (opposite to the large sample sizes available for financial data as in White et al., 2015) suggests considering simplified specifications with limited interaction across quantiles.

Starting from the contribution of Sim and Zhou (2015), we assume that the conditional quantile of a target variable has coefficients dependent both on the chosen quantile of the target variable and on the quantile of one of the covariates. Ideally, the researcher should condition the quantile estimation on the most relevant covariates across the full set of control variables that are available. We identify this most relevant covariate with x_t . The conditional quantile of a target variable is also dependent on the distance for the chosen covariate and its reference quantile. We thus modify Eq. (4) as follows

$$Q_{\tau,\theta}(y_t) = \delta_{0,\tau,\theta} + \delta_{1,\tau,\theta}(x_t - x^\theta) + \delta'_{2,\tau,\theta}Z_t \quad (7)$$

where x^θ is the unconditional quantile of x_t that can be estimated by a standard sample estimator. Note that we can also rearrange terms as

$$Q_{\tau,\theta}(y_t) = \alpha_{\tau,\theta} + \delta_{1,\tau,\theta}x_t + \delta'_{2,\tau,\theta}Z_t \quad (8)$$

where the dependence on the unconditional quantile of x_t is now included in the constant, $\alpha_{\tau,\theta} = \delta_{0,\tau,\theta} - \delta_{1,\tau,\theta}x^\theta$. In order to estimate Eq. (8), we must focus on the observations located in the neighborhood of the θ unconditional quantile of x_t . Again following Sim and Zhou (2015), the estimation of Eq. (8) comes from a minimization problem where the observations are weighted by a Kernel

$$\min_{\alpha_{\tau,\theta}, \delta_{1,\tau,\theta}, \delta'_{2,\tau,\theta}} \sum_{t=1}^T \rho_\tau(y_t - Q_{\tau,\theta}(y_t)) K\left(\frac{F_\theta(x_t) - \theta}{h}\right) \quad (9)$$

where the quantile of y_t is that of Eq. (8), h is a bandwidth, $K(\cdot)$ is a Gaussian kernel that detects the distance between the observed covariate and its quantile by resorting to the distance between the distribution function and the chosen quantile. Note that when estimating the model in Eq. (9), we replace the unconditional distribution function $F_\theta(x_t)$ by a sample estimator. In order to perform the estimation, we consider a plug-in value for the bandwidth set at 0.1, taking into account the limited length of the sample. Given that the model in Eq. (9) requires a sufficiently long sample size to estimate the parameters for a given choice of the quantile of the dependent variable and the relevant explanatory variables' quantile, the research must carefully consider the use of the rolling method.

5. Empirical results

5.1. Linear model

Given the preliminary empirical evidences of Section 3, we start by analyzing the linear dependence of the oil price returns and the changes in the rig count by using the VARX model of Eq. (3). In order to evaluate the potential impacts of the several conditioning variables we introduced in Section 3, we perform a set of preliminary estimations on the full sample (1990-2015), as well as over the ten-year sub-samples (1990-2000, 1995-2005, 2000-2010, and 2005-2015) to detect the most relevant covariates. This full sample division is just a preliminary step to check if the relevant covariates always behave the same under different time periods. Thus, the choice of the sub-samples is merely arbitrary, and is just to make a balance between their sample size (120 data points) and the number of cases to consider. Nevertheless, we emphasize that those subsamples give rise to preliminary estimates needed to verify if the set of covariates impacting either oil returns or changes in rig counts is stable over time or not. The preliminary estimates are available upon requests.

The estimated results suggest that only a subset of the economic and financial covariates introduced in Section 3 is statistically significant in many instances. In this stage, we verified that the oil inventory level, the industrial production, the US dollar index, the VIX, the MOVE and the financial stress indexes, the 10-year Treasury note yields all resulted to be not statistically relevant. On the contrary, the relevant covariates are the change in rig productivity, the change in the effective FED funds rate, which is the arm of conventional monetary policy, the change in the financial condition index, the change in the TED spread which measures changes in liquidity, and the change in the term spread. Notably, these covariates reflect the relevant variables in the oil industry, the conventional monetary policy and financial risk and stress variables.

We thus focus only on the relevant covariates just identified. Moreover, for all of these variables, we include one lag while for the VAR we select three lags. The use of additional lags is not sensibly improving the model fit. We estimate the following VARX model

$$\begin{bmatrix} \Delta Rig_t \\ R_t^{Oil} \end{bmatrix} = \mu + \sum_{j=1}^3 \Phi_j \begin{bmatrix} \Delta Rig_{t-j} \\ R_{t-j}^{Oil} \end{bmatrix} + \beta [\Delta RProd_{t-1} \quad \Delta FFR_{t-1} \quad \Delta FCI_{t-1} \quad \Delta TED_{t-1} \quad \Delta Term_{t-1}] + \varepsilon_t \quad (10)$$

where μ is the 2-component vector of means, each Φ_j is a 2x2 matrix of autoregressive coefficients, and β is a 2x5 matrix of coefficients monitoring the covariates' impacts.

Table 3: Estimation of the VARX model in Eq. (10) over different subsamples

Sample	1990-2015		1990-2000		1995-2005		2000-2010		2005-2015	
Dep. Var.	R_t^{Oil}	ΔRig_t	R_t^{Oil}	ΔRig_t	R_t^{Oil}	ΔRig_t	R_t^{Oil}	ΔRig_t	R_t^{Oil}	ΔRig_t
Expl. Var.										
C	0.005	-0.800	-0.001	-3.719	0.015	-1.641	0.012	1.695	0.002	0.521
	1.050	-0.488	-0.157	-1.509	1.786	-0.953	1.412	0.957	0.205	0.174
R_{t-1}^{Oil}	0.052	84.111	-0.073	43.293	-0.133	21.470	0.081	38.731	0.199	100.255
	0.874	4.389	-0.817	1.420	-1.406	1.137	0.876	2.085	1.931	2.689
R_{t-2}^{Oil}	-0.006	81.605	-0.064	49.716	-0.160	45.056	0.061	54.691	0.098	92.165
	-0.098	4.186	-0.699	1.599	-1.672	2.345	0.684	3.058	0.946	2.465
R_{t-3}^{Oil}	0.131	53.766	0.230	67.250	0.026	36.228	0.056	32.564	0.137	67.517
	2.244	2.838	2.772	2.368	0.263	1.830	0.621	1.785	1.237	1.689
ΔRig_{t-1}	0.000	0.517	-0.001	0.014	0.000	-0.019	0.001	0.232	0.000	0.594
	-0.157	8.049	-1.853	0.110	0.005	-0.132	2.147	1.919	0.293	5.948
ΔRig_{t-2}	0.000	0.022	0.000	-0.143	0.000	-0.151	-0.001	0.091	0.000	0.035
	-0.355	0.356	-0.267	-1.526	-0.110	-1.561	-2.507	0.979	-0.647	0.329
ΔRig_{t-3}	0.000	0.101	0.000	-0.048	0.000	0.034	0.000	0.230	0.000	0.122
	-1.468	1.812	-1.633	-0.534	-0.638	0.375	-1.105	2.632	-0.360	1.332
$\Delta RProd_{t-1}$	0.006	3.139	0.003	0.412	0.004	-0.286	0.009	0.698	0.012	5.365
	2.287	3.895	0.877	0.335	0.888	-0.349	1.946	0.748	1.644	1.957
ΔFFR_{t-1}	0.123	-8.707	0.212	12.227	0.201	13.450	0.046	6.795	0.008	-16.374
	3.450	-0.754	4.367	0.736	3.099	1.037	0.815	0.598	0.130	-0.701
ΔFCI_{t-1}	0.018	4.835	-0.011	2.685	0.009	2.007	0.053	3.530	0.026	5.254
	1.759	1.437	-0.638	0.448	0.424	0.475	3.846	1.261	1.778	0.990
ΔTED_{t-1}	-0.021	13.611	-0.085	38.920	-0.046	11.103	0.081	9.226	0.008	14.679
	-0.665	1.325	-1.347	1.812	-0.620	0.744	1.985	1.126	0.202	0.993
$\Delta Term_{t-1}$	0.040	-7.027	0.092	-15.149	0.063	-1.147	-0.022	6.152	-0.012	-0.956
	1.733	-0.940	2.814	-1.353	1.700	-0.154	-0.604	0.823	-0.272	-0.062

<i>AdjR</i> ²	0.094	0.423	0.204	0.109	0.041	0.055	0.160	0.291	0.112	0.583
--------------------------	-------	-------	-------	-------	-------	-------	-------	-------	-------	-------

Notes: The bold values denote the 10% statistically significant coefficients and the associated t-statistics. R_t^{oil} is the return on oil, ΔRig_t is the change in rig counts, C is the constant, $\Delta RProd_{t-1}$ is the change in rig productivity, ΔFFR_{t-1} is the change in the federal funds rates, ΔFCI_{t-1} is the change in the financial condition index, ΔTED_{t-1} is the change in the TED spread, $\Delta Term_{t-1}$ is the change in the term spread. The standard errors have been estimated using a Newey-West estimator for the covariance of the residuals. The last line reports the adjusted R-squared for each of the two equations composing the VARX model.

From Table 3, we note that the lags of the oil returns are very relevant in explaining the dynamics of the changes in the rig count, both in the full sample as well as in the most recent 10-year sub-samples. In all cases, the coefficients are positive meaning that an increase (a decrease) in the oil prices over the previous quarter leads to an increase (a decrease) in the rig counts of the current month. The larger the change in the oil price, the larger the effect on the active rigs, underscoring the importance of the size of oil revenues on drilling activity. This implies that the recent 60% drop in oil prices should have a strong negative effect on oil rig counts, with the possibility of a balancing act on the oil glut. On the contrary, lags of the changes in rig counts have only minor effects on the dynamics of changes in rig counts. Nevertheless, we note that there is a somewhat stable impact of the lagged rig count change on the rig count change. This effect is stronger in the most recent years. This might suggest that the series dynamics are being characterized by a more persistent behavior in the recent years.

The impact of changes in rig count on oil returns is limited and appears only in the recent years. This highlights the importance of the recent periods, which include the introduction of shale fracking. In recent years, shale oil production increased by more than 40% and this increase comes, in the largest part, from thousands of small companies that sought to rapidly increase production to pay off their loans. When an oil glut developed in the market and prices plunged, market participants and the media started to pay increasingly more attention to the rig accounts than in previous periods, aided by the fact that oil companies started to curtail their

capital expenditures and reduced drilling activity. Moreover, we emphasize that the impacts of the other covariates, the FED fund rates, the Financial Condition Index, the TED and Term spreads, monitoring the overall economic and financial conditions play only a minor role.

To evaluate the variation over time in the model fit, we look at the time evolutions of the two R-squareds for the oil returns and the changes in the rig count. We note the relevant changes in those adjusted R-squareds of both equations as both sensibly oscillate over the periods. In fact, while the full-sample's R-squared result for the oil returns is around 10%, the over sub-samples' results range from 4% in the 1995-2005 subsample to 20% in the 1990-2000 subsample, and in the most recent subsample last period the result is around 11%. On the other hand, the R-squared for the changes in the rig count varies from 5% in 1995-2005 to a peak of 58% in 2005-2015. Those movements suggest that the structural relationships across the variables may be changing over time, either for changing economic conditions (going thus beyond to what is captured by the economics and financial variables we consider) or because of the presence of a structural break as a result of the shale fracking as discussed earlier. More possibly, there is an increase in the responsiveness of rig counts to changes in oil prices facilitated by shale oil fracking.

As a further analysis, we report in Figure 4 the rolling estimation of the two adjusted R-squareds (for oil returns and changes in rig counts) over a 5-year estimation window. Few elements clearly appear. First, up to 1998, the two R-squareds move together following a decreasing trend. Later, from 1999 up to 2005, they follow opposite directions: the oil returns' R-squared shows first an increase followed by a decrease. On the other hand, the change in rig counts' R-squared experiences first an increase and then a decrease. Both R-squareds move together in the last part of the sample, from 2006 to 2015 with common drops in the model fit

after the global financial crises and then again in 2013. A quick recovery follows both drops. We attribute the most recent upward trend in the model fit to the production coming from shale oil, in particular from the thousand small sized shale companies.

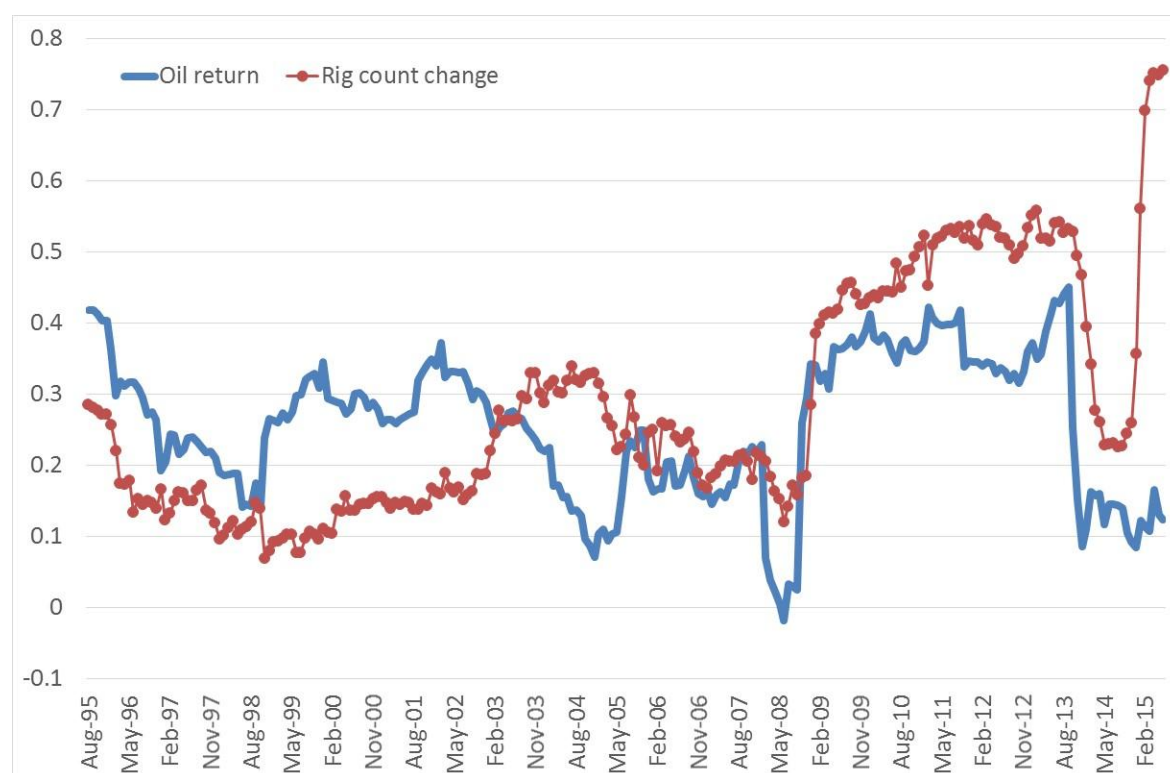


Figure 4: adjusted R-squared evolutions over a 5-year estimation window for the model in Eq. (8).

Notes: The adjusted R-squared for the oil return is the solid blue line, while that for the change in rig count is the circled red line.

The change over time in the R-squareds further confirms the possible presence of a structural break. We thus perform a standard structural break test, the Chow test, focusing on the entire model, or on the two separate equations forming the model. Given that, we do not have a-priori knowledge of the break date location, we thus run the Chow test moving the break date from 1995 to 2005 and testing for break on each month. The null hypothesis is that in all cases

the parameters before and after the break data are stable. Figure 5 reports the p-values of the three tests (break in the full model, break in oil equation, and break in rig count equation). If we focus on the full model, we note that the p-values are constantly below the 1% level, suggesting a high instability of the parameters. However, by looking at the single equation outcomes, we note two opposite patterns. First, for the oil return equation, and taking the 1% confidence level for rejection of the null, we detect a break only up to 2002. As from 2003, the p-values stay above 0.01, with a temporary decrease during the global financial crisis period. Stating it differently, for the change in the rig count, we start detecting a break only from 2003 onward. The oil prices began to move up significantly in 2003 starting with the 2003 Iraq war, and then continued through the commodity boom that dominated the period 2004-2006.

Taking into account the limited dynamic impact of the changes in the rig counts on the oil returns, the different dynamic behavior of the two variables, and the fact that, to our best knowledge, is the first contribution that analyzes the changes in rig count series with linear and non-linear models, we from now on focus only on the changes in rig counts, leaving aside further analyses on the oil return series.

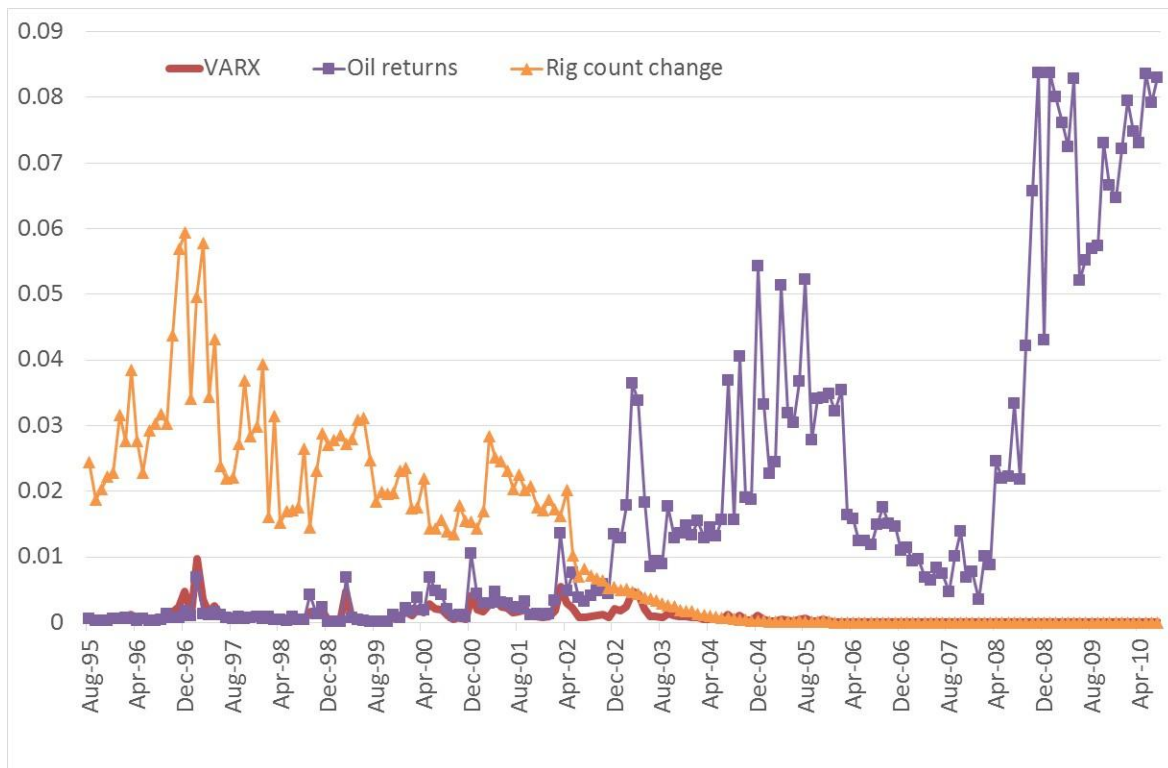


Figure 5: The p-values for the Chow break test for different break dates for the VARX model.

Notes: The trajectories correspond to the p-values for the null hypothesis of no break on the model, or on one single equation of the model. The blue trajectory with squares refers to the p-value of the Chow test for the oil returns, the orange with triangles represents the trajectory for the p-values for the Chow test on changes in rig counts, and the red solid line trajectory denotes the p-values for the Chow test conducted on the VARX model that contains the two equations.

5.2. Quantile models

We now consider the analysis of the relationship between the changes in rig counts and the oil returns by means of quantile methods, thus underscoring the potential of non-linearity and asymmetry across the quantiles of the variables. Given the discussion and evidences of the previous section, when moving to the analysis of the quantile regression models, we focus only on the specifications where the change in rig counts is the dependent variable. Moreover, as we plan to consider the possible non-linear impact from oil returns on the changes in rig counts, as modelled by means of the quantile regression and the quantile-on-quantile analyses, we allow for

two possible lag designs. In the first case, we include among the quantile covariates three monthly lags of oil returns, while in the second we consider the previous quarter oil return lag only. We also include in both cases the last period's lag for the change in oil productivity as a further empirically relevant covariate. We exclude all the additional covariates adopted in the linear model of Eq. (10) as they turned out to be of a limited statistical significance over the full sample.⁸

We thus consider the following models for the conditional quantiles

$$Q_{\tau}(\Delta Rig_t) = \delta_{0,\tau} + \delta_{1,\tau}\Delta Rig_t + \delta_{2,\tau}R_{t-1}^{Oil} + \delta_{3,\tau}R_{t-2}^{Oil} + \delta_{4,\tau}R_{t-3}^{Oil} + \delta_{5,\tau}\Delta RProd_{t-1} \quad (11)$$

$$Q_{\tau}(\Delta Rig_t) = \delta_{0,\tau} + \delta_{1,\tau}\Delta Rig_t + \delta_{2,\tau}(R_{t-1}^{Oil} + R_{t-2}^{Oil} + R_{t-3}^{Oil}) + \delta_{3,\tau}\Delta RProd_{t-1} \quad (12)$$

The use of a single covariate associated with the oil returns also allows for using the specification of a quantile-on-quantile model

$$Q_{\tau,\theta}(\Delta Rig_t) = \alpha_{\tau,\theta} + \delta_{1,\tau,\theta}\Delta Rig_t + \delta_{2,\tau,\theta}(R_{t-1}^{Oil} + R_{t-2}^{Oil} + R_{t-3}^{Oil}) + \delta_{3,\tau,\theta}\Delta RProd_{t-1} \quad (13)$$

where the first quantile index, τ , refers to the change in rig counts, while the second quantile, θ , points to the quarterly oil return.

We emphasize that the quantile models exclude all the covariates previously analyzed in the linear model with the exception of the rig productivity, which is the most relevant covariate for changes in the rig counts. With too many parameters to estimate (that is with many

⁸ This decision is based only on the statistical significance, which is limited for all covariates we included in the VARX model and whose estimation results are reported in Table 3. Moreover, note that the estimation of the quantile regression requires either long samples or a limited number of coefficients, as we are slicing the density support of the dependent variable.

covariates), the efficiency of the estimators gets lower and we might end up with no significance at all. Further, we do not estimate the previous models on the full sample as even the simple linear specification shows parameter instability. We thus start directly with the rolling evaluation of the quantile regression models in Eqs. (11) and (12).

Notably, when estimating the quantile regression models, we recover the coefficients that are quantile-specific. Therefore, despite Eqs. (11) and (12) having a limited number of parameters to estimate, the total number across quantiles and over a rolling scheme is too large to be included in a table. We thus decide to provide a first graphical evaluation monitoring the statistical significance of the estimated coefficients. We fix the quantiles we consider in the range 0.1 to 0.9 with a 0.1 step. Moreover, we fix the estimation window to 120 data points (ten years as we use monthly frequency), and we roll the sample with a 1-month step. Reminding that the full sample has a size of 298 months, we have a sequence of 179 estimates of the models in Eqs. (11) and (12), each is based on a different 120-month window. For each of those 179 estimates of the models, we count across quantiles the fraction of statistically significant coefficients associated with the change in rig counts, oil returns and the change in rig productivity. If a covariate is always statistically significant in causing the change in rig count quantiles (across the 9 quantiles we consider), the plot will reach the 100% level.




Figure 6: Fraction of statistically significant coefficients across quantiles and over time for the three different covariates of the model in Eq. (11).

Notes: The blue color refers to the fraction for the oil returns, the red points to the fraction for the change in rig counts, and the green color denotes the fraction for the change in rig productivity.




Figure 7: Fraction of statistically significant coefficients across quantiles and over time for the three different covariates of the model in Eq. (12).

Notes: See the notes of Figure 6 for more information.

Figures (6) and (7) provide interesting insights into the impact of the selected covariates on the change in the rig count quantiles. First, for the lagged change in rig counts and the lagged change in the rig productivity, the use of three monthly lags of oil returns or the use of the last quarter's oil return is almost irrelevant. In both figures, we note that the two covariates are statistically significant only in the second half of the full sample that is from 2010 onwards. This is consistent with the results of the linear model where the two variables become statistically significant only from the sample 2000-2010. Note that in Figures 6 and 7, the dates refer to the end of the rolling sample, thus 2010 is associated with the samples starting in 2000). Further, the

impact of the changes in rig counts is far more relevant than the impact of changes in the rig count productivity.

Moving to the oil returns, we see a marked difference between the two plots. In fact, on the one side, the use of three monthly lags shows evidences of periods where there seems to be no causality from oil returns to the changes in rig counts (for instance for the 10-year samples ending in 2002), or a very limited statistical significance (samples ending in 2014). On the other side, the use of the previous quarter returns provides a more stable pattern, with a limited impact for the 10-year samples ending up to 2002-2003, and then a marked increase, peaking even at 100%, with a large significance up to the end of the sample. We emphasize that the obvious serial correlation of the quarterly oil returns (which are computed on a monthly basis using the last three months) in the model defined by Eq. (10) does not impact the changes in rig counts' own dynamics which remain significant in the last part of the sample.




Figure 8: Time-variation of the quantile regression coefficients for quarterly oil returns.

Notes: The model is that of Eq. (12) evaluated for different quantiles. The solid line is the 90% quantile and the dashed line is the 10% quantile. Other quantiles assume in most cases values between or close to those two limiting quantiles.

Given the previous observations, the model with the quarterly oil returns seems more stable than the specification given in Eq. (11). Focusing thus on the model of Eq. (12), we may get further insights by looking at the time variation in the coefficients over time of the impact of the oil returns. This is plotted in Figure 8, where we highlight the two extreme quantiles' coefficients (10% and 90%). We note that, in the first part of the figure, for the 10-year samples ending up to 2007, the coefficients are quite different across quantiles. In fact, they move from very large values on the left quantiles (below median or bearish markets) to small or even negative (but not statistically significant) values for the right quantiles (above median or bullish markets). In the second half of the sample, the coefficients stabilize and are closer one to the

other, taking values in the range of 20-60. In this latter period, increases in the oil returns move to the right the density of the changes in rig counts, affecting only the location and not the shape. Therefore, it seems that the relationship between the two variables is moving toward linearity in the most recent years. This calls for a proper test to evaluate the stability of the quantile regression coefficients across quantiles and over time. In fact, if the coefficients are constant across quantiles, we might question the appropriateness of using the quantile regression framework, thus implicitly validating or strengthening the use of the linear models.

Figure 9 reports the test for the equality of the quantile regression coefficients on each of the rolling estimates. The null hypothesis of the test is that all the coefficients (or just the coefficients of selected covariates), excluding the constant, are equal across quantiles. The p-values clearly show that the coefficients are stable across the quantiles in the second part of the sample, while some mild evidences of differences across the quantiles exist (in particular for the oil coefficients) in the first part of the sample. Nevertheless, we stress that the test power, or the quantile regression estimates, might be affected by the limited sample size adopted in the rolling estimation approach.

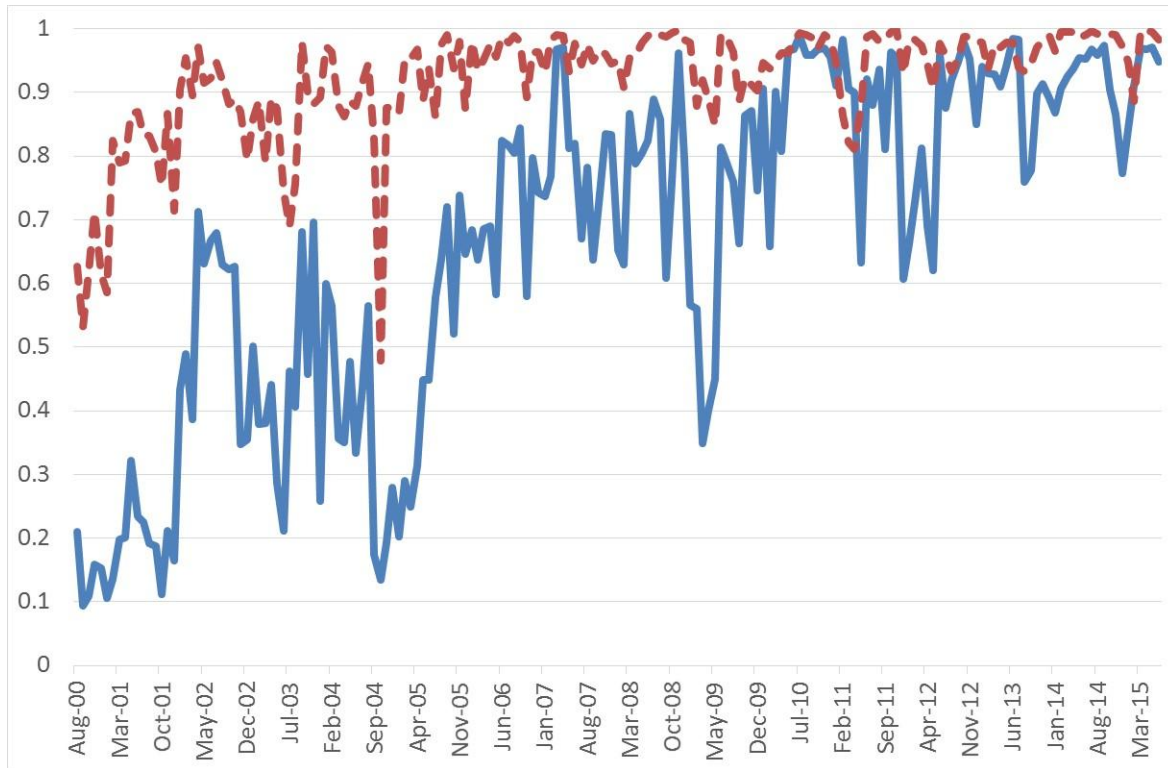


Figure 9: The p-value (upper tail) of the test for the equality across quantiles.

Notes: The test for all coefficients in the model of Eq. (12) is the solid blue line. In dashed red line, the same test but is restricted to the equality of the coefficients of the quarterly oil returns.

Even if the quantile regression framework of Eqs. (11) and (12) seems to provide only weak evidences of the change in the impact of oil returns on the changes in rig counts across the quantiles, it confirms that the relationship across the two variables in the first part of the sample is different from that in the second part of the sample. In the most recent years, the impact of oil returns is more stable and there is a relevant role for the lagged changes in rig counts and the lagged change in rig productivity. The impact for those two variables when they are statistically significant is always positive, suggesting that an increase (a decrease) in the lagged rig counts or the lagged rig productivity moves upward (downward) the density of the rig count changes. Therefore, an increase in the rig productivity or an increase in the rig count at time t , leads to an

increase in the probability of observing an increase in the rig counts at time $t+1$. The detailed results for all coefficients are available upon request.

As a further and final empirical check, we focus on the quantile-on-quantile approach of Eq. (13). Given the evidences of the break between the first part and the second part of the full sample, and given the need for larger sample sizes to estimate the model of Eq. (13), we perform just two estimates on the two parts of the full sample, and contrast them with the full-sample analysis. Note that the number of estimated coefficients is now increasing as they are indexed on both the change in the rig count quantiles and on the oil return quantiles. Thus, while still focusing on a 10% grid over the quantiles, we do have 81 estimates of Eq. (13) for each considered sample. Similarly to Sim and Zhou (2015), we graphically represent the estimated coefficients by means of surface plots, see Figures (10) to (12). In our case, the most relevant impact is that of the quarterly oil return, and thus we provide plots only for this variable. Similar plots for the other covariates are available upon request.

The three figures adopt the same scale and color grading so that a comparison is immediate. We clearly see a change between the 1997-2002 period and the 2003-2015 subsample. In the first period, the impact of the oil returns on the rig counts is high when the oil returns are in their lower quantiles (bearish markets), but from the median upward, the impact on the counts' quantiles stabilizes and is no longer changing over the oil return quantiles. In the most recent years, the oil return impact is not affected by the oil return quantiles and is much more stable, probably due to the increasing glut in the oil market. We read this as further evidence of the stability of the relationships across the variables, consistently with the outcome of the quantile regression estimates.




Figure 10: The full-sample quantile-on-quantile estimates of the impact of the previous quarter oil returns on the change in rig counts by conditioning on the corresponding quantiles.

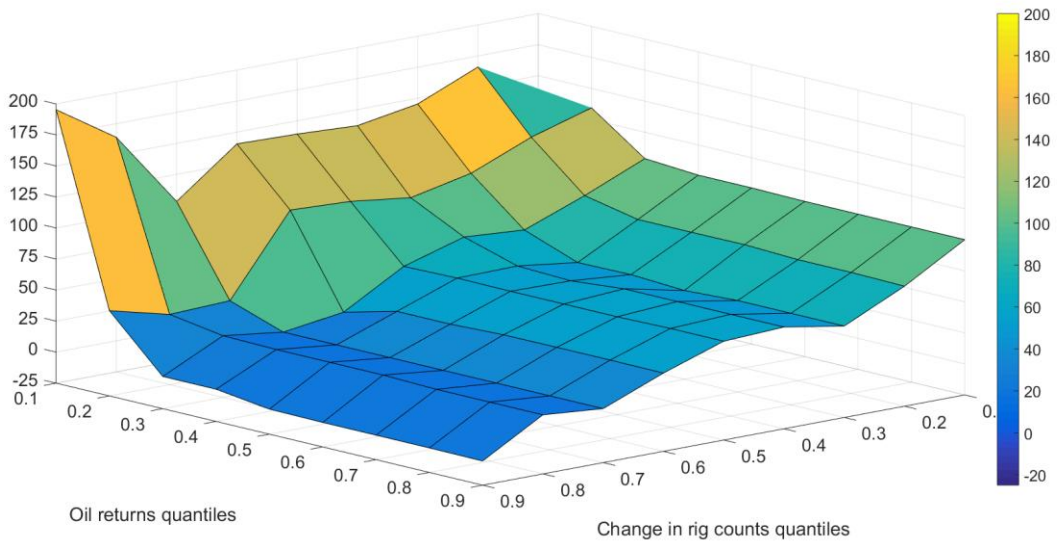


Figure 11: The 1990-2002 quantile-on-quantile estimates of the impact of the previous quarter oil returns on the change in rig counts by conditioning on the corresponding quantiles.




Figure 12: 2003-2015 quantile-on-quantile estimates of the previous quarter oil returns impact on the change in rig counts by conditioning on the corresponding quantiles.

6. Conclusion

The study deals with a timely and relevant issue in the oil markets which is the relationship between changes in oil prices and changes in rig counts, while reckoning for control variables. This relationship is of significant interest to analysts, investors, oil companies, commercial banks, investment banks and investors among others. The descriptive analyses is consistent with the previous literature and provides evidence of the presence of positive lagged relationships between oil returns and changes in rig counts, which is predominantly strong when the impact is from changes in oil prices to changes in rig counts. Moreover, we do have evidence suggesting the presence of a non-linear link between the variables. This is particularly clear if we use exceedance correlations. Nevertheless, we first use the linear analysis to explore this relationship between the oil returns and changes in rig counts and find instabilities in the relationship as supported by the break test. Therefore, we move to non-linear methodologies.

We use the quantile regression and quantile-on-quantile analyses because this framework is flexible and offers a comprehensive approach for examining how the explanatory variables affect the location, scale and shape of the entire change in the rig count distribution. The two approaches are especially significant in estimating relationships that change across the quantiles, and thus detecting asymmetric interdependencies among the variables.

The study provides several important outcomes. All the analyses point out that there exists a positive relationship between changes in oil prices and changes in oil counts but that this relationship is not contemporaneous. In fact, we observe lagged oil returns positively affecting the change in rig counts, with a lag going up to one quarter. This implies that changes in oil prices are not immediately reflected in changes in the rig count. Further, if we restrict the attention to linear models, we note that the relevance of the relationship changes over time but becomes stronger and more stable from 2005 onward, even with controlling for the impacts of other covariates and the lagged change on rig counts. Break tests confirm the presence of this change in the relationship that in the most recent years becomes steadier.

When moving to non-linear models, the quantile regression outcomes confirm the importance of lagged oil returns in the causation of change in rig count quantiles, with increased statistical significance from 2005 onward. However, the impact seems to be constant across the quantiles, thus suggesting the appropriateness of linear specifications. To better understand the nature of the information flow from oil returns to change in rig counts, we fit a quantile-on-quantile regression, where we condition coefficients on quantiles of both oil returns and the change in rig counts. We find that the relationship between the variables is stronger and less stable in the below the median quantiles of change in rig counts. Moreover, we note that the impact of oil returns on the changes in rig counts becomes higher when the oil returns take on

values in the lower quantiles. Therefore, oil impact is stronger during bearish oil markets. This result may have an implication for the recent 70% plunge in the oil prices. It implies that the ensuing fall in oil rig counts will be large, but it will come with a lag. In turn, this has strong consequences for the prevailing oil glut over time. Moreover, since the literature shows that changes in oil rig count affect production, then this serial result is relevant for the duration of the oil glut as it must be shorter than the market expects, or it may point to a significant drop in U.S. oil production but also with a lag.

However, we also note that the impact is stabilizing in the recent years, moving the relationship across variables toward linearity, even if with the presence of lags. A linear relationship, thanks to its simple structure, allows for immediate analyses coming from changes in the covariates, which in this case are the oil returns. Therefore, irrespective of the size of the changes in rig counts or of oil returns, we are able to evaluate the consequent future evolution of changes in rig counts, thanks to the presence of lags. Finally, we stress that in the above median changes in rig count quantiles, the relationship is stable and positive, implying that bullish oil markets increase the rig counts. It is possible the shale oil boom is a confounding factor in the relationship, which shows significant structural breaks.

References

Abraham, K. S., 2000. A revised global drilling outlook for 2000 and update on Iraq. *Journal of Energy and Development* 25, 203-215.

Anderson, S. T., Kellogg, R., and Salant, S. W. (2014). *Hotelling under pressure*. Energy Institute at Haas, Working paper WP 250.

Black, G., LaFrance, J. T., 1998. Is Hotelling's rule relevant to domestic oil production? *Journal of Environmental Economics and Management* 36, 149-169.

Brown, J. P. (2015). The response of employment to changes in oil and gas exploration and drilling. *Economic Review*, second quarter, 57-81.

Ghoury, S. and Aneesuddin, M. (2015). Cyclical oil prices – Is it a necessary condition to balance global oil supply and demand? TDJ World News. <http://thedailyjournalist.com/the-financier/cyclical-oil-prices-is-it-a-necessary-condition-to-balance-global-oil-supply-and-demand/>

Hunt, L.C., Ninomiya, Y., 2003. Unravelling trends and seasonality: a structural time series analysis of transport oil demand in the UK and Japan. *Energy Journal* 24 (3), 63-96.

Kellogg, R. (2014). The effect of uncertainty on investment: Evidence from Texas oil drilling, *American Economic Review*, 104(6), 1698–1734.

Koenker, R., 2005, *Quantile Regression*, Econometrics Society Monographs n. 38, Cambridge University Press.

Koenker, R., and Basset, G., 1978, Regression quantiles, *Econometrica*, 46(1), 33-50.

Longin, F.M. and B. Solnik (2001), Extreme correlation of international equity markets, *Journal of Finance*, 56, 649-676.

Majeure, F. (2015). Crude Oil: The divergence of rig count, oil price and production Explained. Seeking Alpha.com.

Mohn, K. and Osmundsen. P. (2008) Exploration economics in a regulated petroleum province: the case of the Norwegian Continental Shelf. *Energy Economics* 30 (2), 303-320.

OGJ, (2003). Increased drilling spurs industry towards recovery. *Oil and Gas Journal*, Sept. 22, 49-61.

Osmundsen, P., Sørenes, T., Toft, A. (2008). Drilling contracts and incentives, *Energy Policy* 36, 3138– 3144.

Osmundsen, P. and Mohn, M. (2006) Valuation of international oil companies, *The Energy Journal*, 49-64

Ringlund, G. B., Rosendahl, K. E. and Skjerpen, T. (2008). Does oilrig activity react to oil price changes? An empirical investigation. *Energy Economics* 30 (2), 371- 396.

Toews, G and Naumov, A. (2015). The relationship between oil price and costs in the oil and gas industry. *Oxford Centre for the Analysis of Resource Rich Economies*, 152, University of Oxford, Oxford, U.K.

White, H., Kim, T., and Manganelli, S., 2015, VAR for VaR: measuring tail dependence using multivariate regression quantile, *Journal of Econometrics*, 187, 169-188.