

Has Oil Price Predicted Stock Returns for Over a Century?#

Paresh K. Narayan* and Rangan Gupta**

ABSTRACT

This paper contributes to the debate on the role of oil prices in predicting stock returns. The novelty of the paper is that it considers monthly time-series historical data that span over 150 years (1859:10-2013:12) and applies a predictive regression model that accommodates three salient features of the data, namely, a persistent and endogenous oil price, and model heteroskedasticity. Three key findings are unraveled: First, oil price predicts US stock returns. Second, in-sample evidence is corroborated by out-sample evidence of predictability. Third, both positive and negative oil price changes are important predictors of US stock returns, with negative changes relatively more important. Our results are robust to the use of different estimators and choice of in-sample periods.

JEL classification: C22, E37, G17, Q43

Keywords: Stock returns, Predictability, Oil price

We would like to thank two anonymous referees for many helpful comments. However, any remaining errors are solely ours.

*Centre for Financial Econometrics, School of Accounting, Economics and Finance, Deakin University, Australia. Corresponding author. Email: paresh.narayan@deakin.edu.au.

** Department of Economics, University of Pretoria, Pretoria, 0002, South Africa. Email: rangan.gupta@up.ac.za.

I. Introduction

The impact of oil price on stock returns have now become a familiar subject of research in financial and energy economics (see, for instance, Driesprong 2008; Narayan and Sharma, 2011). The literature has explained the effect of oil price on stock returns through the gradual diffusion hypothesis, proposed in the asset pricing literature by Hong and Stein (1999) and Hong *et al.* (2007), which perceives that stock returns underreact to oil price news.

In this paper, we re-visit the relationship between oil prices and stock returns. Our empirical investigation is based on four approaches, each different from the literature. First, we use the longest historical time-series data possible. Our data set spans over 150 years of monthly data (1859:10-2013:12), covering the entire modern era of the petroleum industry, which typically began with the first drilled oil well in Titusville, Pennsylvania in 1859. This unique data set allows us to understand the oil price—stock returns relationship from a historical perspective. In this regard, our paper represents the first historical analysis that makes use of data dating to 1859.

Second, we examine the oil price—stock returns relationship using a time-series predictive regression model based on a generalized least squares estimator recently proposed by Westerlund and Narayan (2012, 2014). The attractiveness of this test (estimator) lies in its ability to accommodate both a persistent and endogenous oil price variable, and any heteroskedasticity in the regression model. These statistical features are accepted as a stylized fact when one is using daily financial time-series data, therefore, appropriately treating them is imperative.

Third, unlike the extant literature on stock prices (returns) and oil prices which have relied on in-sample inference, we conduct our analysis using both in-sample and out-of-sample empirical frameworks. This is of paramount importance, since existence of in-sample predictability does not necessarily ensure out-of-sample forecasting gains; see, for example,

Rapach and Zhou (2013). Campbell (2008: p.3) stresses this point succinctly, when he suggests that: “The ultimate test of any predictive model is its out-of-sample performance”.¹

Finally, we examine whether oil price nonlinearly predicts US stock returns. Specifically, we use the Granger and Yoon (2002) approach to filter oil price series into a positive and a negative oil price series. This allows us to test whether positive and negative oil prices predict US stock returns any differently. A nonlinear test is important because in a recent study, Narayan and Sharma (2011) show that oil price has a nonlinear effect on sectoral stock returns for stocks listed on the New York Stock Exchange. By addressing any potential nonlinear relationships over a historical time period of data brings completeness to our empirical analysis.

Our approaches lead to the following key findings. Our first finding is that oil price is a persistent and endogenous predictor variable and that our proposed stock return predictability model is heteroskedastic. Second, using an estimator that accounts for these three salient features of the data ensures that we reject the null hypothesis of no stock return predictability. Our third finding is that while negative and positive oil prices predict US stock returns, this predictability is indeed nonlinear, in that, negative oil prices predict US stock returns more than do positive oil prices.

Finally, we undertake a range of robustness tests. We do this as follows. First, we use a range of out-of-sample forecasting evaluation metrics and confirm that the in-sample evidence of predictability is corroborated by out-of-sample tests. In other words, our proposed predictive regression model always beats the constant returns model. Second, an issue at the heart of out-of-sample forecasting evaluations, which has implications for robustness test outcomes, is the

¹ It is important to point out that there is a pendulum of arguments supporting and similarly chastising out-of-sample tests, just like in-sample tests (Narayan et al., 2014). There is no theory to guide applied researchers on this. We do not want to be caught-up in this debate. Therefore, we do both tests. We see this as an important step because it serves as a robustness measure.

choice of the in-sample period. Typically, in the absence of any theoretical guidance on this, researchers have used in-sample periods of 25%, 50% and 75%. Westerlund and Narayan (2012) argue that this treatment is sufficient for a robustness test of out-of-sample performance. Our empirical analysis is free of criticisms labelled at the choice of the in-sample periods because we, following Westerlund and Narayan (2012), consider all three in-sample periods. Doing so, we conclude that results render our predictive regression model superior to the constant returns model, regardless of the choice of the in-sample period.

Our approaches and findings contribute to two specific strands of the literature. Our first contribution is to the literature which has shown that oil prices predict/impact stock returns. Gupta and Modise (2013), for instance, show that different oil price shocks affect South African stock returns differently; that oil prices matter for European and emerging stock returns have been confirmed by Cunado and de Gracia (2014) and Asteriou and Bashmakova (2013), respectively; the effect of oil price shocks on industry/sector returns have been documented in Elyasiani et al. (2011) and Lee et al. (2012); and, there are several individual country-based studies that confirm a statistical relationship between oil prices and stock returns (see Ghosh and Kanjilal, 2014 and the references therein). Typically these studies, discussed above, use 30-50 years of data to show that oil prices predict stock returns. We show that this is true even when one uses 150 years of data. The main message we provide is that the relationship between stock returns and oil price is not something new; it is a predictability relationship which spans over one-and-a-half century. The implication is that oil price and stock markets share a long-run relationship, an outcome that will be of interest for empirical models of forecasting and for theoretical models that require as a pre-requisite a cointegrating relationship between oil price and the stock market.

Recent studies also show that the relationship between oil price and stock returns are nonlinear. Narayan and Sharma (2011) show that US stock returns and oil price are nonlinear in five out of 10 sectors making up the NYSE. Similar findings have been reported by Lee and Zeng (2011) using market-level data. We demonstrate nonlinearity in this relationship through using a different approach. Specifically, using the Granger and Yoon (2002) approach we extract positive and negative oil price changes and use these time-series to test for stock return predictability. By showing that both positive and negative oil prices predict returns with negative oil price changes moving returns most, we contribute to the robustness of the evidence that suggests that oil prices predict US stock returns in a nonlinear fashion. The implication is that future studies on this subject should consider nonlinear models in addition to using linear models.

The balance of the paper appears as follows. The next section discusses the estimation approach. The data and empirical results appear in Section 3. The final section concludes with the key findings.

2. Estimation Approach

2.1. Time-series predictive regression model

Following the literature, we use the following bivariate predictive regression model:

$$y_t = \alpha + \beta OP_{t-1} + \varepsilon_{y,t}. \quad (1)$$

Here, y_t is the stock returns, computed as log difference, in month t and OP_t is the oil price in the same month. The null hypothesis of no predictability is $H_0: \beta = 0$. As explained earlier, in the above specification, it is possible that oil price is endogenous and persistent. If it is,

one can expect a bias, leading to deceptive inference on the no predictability null. To avoid this, we follow Westerlund and Narayan (2012, 2014) and model oil price as follows:

$$OP_t = \mu(1 - \lambda) + \lambda OP_{t-1} + \varepsilon_{OP,t} \quad (2)$$

where $\varepsilon_{OP,t}$ is mean zero and with variance σ_{OP}^2 . If the error terms from Equations (1) and (2) are correlated, then oil price is perceived to be endogenous. In order to allow for this possibility, we assume that the error terms are linearly related in the following way:

$$\varepsilon_{y,t} = \theta \varepsilon_{OP,t} + \epsilon_t \quad (3)$$

where ϵ_t is again mean zero and with variance σ_{ϵ}^2 .

We use two estimators, namely, the bias-adjusted OLS estimator proposed by Lewellen (2004) and GLS estimator proposed by Westerlund and Narayan (2012, 2014). Both estimators are based on making Equation (1) conditional on Equation (2), thereby, removing the effect of the endogeneity and accounting for any persistence in the predictor variable. The resulting conditional predictive regression model can be written as follow:

$$y_t = \alpha + \theta\mu(1 - \lambda) + \beta^{adj} OP_{t-1} + \theta OP_t + \epsilon_t \quad (4)$$

where ϵ_t is independent of $\varepsilon_{OP,t}$ by construction and $\beta^{adj} = \beta - \theta(\lambda - 1)$. The bias-adjusted OLS estimator of Lewellen (2004) is basically the OLS estimator of $\beta^{adj} = \beta - \theta(\lambda - 1)$ in Equation (4). The key difference between this estimator and the one of Westerlund and Narayan (2012, 2014) is the accounting for potential conditional heteroskedasticity in ϵ_t . Lewellen (2004) uses OLS, which means that any information contained in the heteroskedasticity is ignored. The GLS estimator, on the other hand, exploits this information and is, therefore, expected to be more precise. In particular, it is assumed that ϵ_t has an autoregressive conditional heteroskedastic (ARCH) structure. Westerlund and Narayan (2012) suggest that if one uses a higher order ARCH, which can be mimicked in-sample by fitting an AR model to the squared OLS residuals

obtained from equation (3), $\hat{\epsilon}_t$ say. The fitted value from this AR model, $\hat{\sigma}_{\hat{\epsilon}_t}^2 = \hat{\psi}_0 + \sum_{j=1}^q \hat{\psi}_j \hat{\epsilon}_{t-j}^2$, is a consistent estimator of $\sigma_{\hat{\epsilon}_t}^2$ and can, therefore, be used as a weight when performing GLS. The Westerlund and Narayan (2012) test for predictability (or rather the absence thereof) is the resulting GLS t -statistic for testing $H_0: \beta^{adj} = 0$ in Equation (4).

2.2. Approach to extracting negative and positive oil price series

As indicated above, since the literature shows a nonlinear relationship between oil price and stock returns, entertaining this type of nonlinearity in a predictive regression framework is imperative. Our approach is to deal with a nonlinear oil price by disaggregating it into positive and negative price components. This idea fits in with the proposal for disaggregating a time-series into positive and negative changes recommended by Granger and Yoon (2002). Using their approach, we begin by defining oil price as a random walk:

$$OP_t = OP_{t-1} + u_{1t} = OP_0 + \sum_{i=1}^t u_{1i} = OP_0 + \sum_{i=1}^t u_{1i}^+ + \sum_{i=1}^t u_{1i}^-, \quad (5)$$

where $t = 1, 2, \dots, T$. OP_0 is a constant representing the initial value of oil price, u_{1i} indicates a white noise error term, which is defined as the sum of positive and negative shocks, i.e., $u_{1i} = u_{1i}^+ + u_{1i}^-$, where $u_{1i}^+ = \max(u_{1i}, 0)$ and $u_{1i}^- = \min(u_{1i}, 0)$. The cumulative form of OP when it is positive is: $OP^+ = \sum_{i=1}^t u_{1i}^+$, and when OP is negative, it is: $OP^- = \sum_{i=1}^t u_{1i}^-$. Now we have two additional time-series of oil price variables, namely, a negative oil price series and a positive oil price series. The advantage is that we are not only able to test whether oil price predicts US stock returns, but also whether negative and positive changes in oil prices predict stock returns.

3. Data and Empirical Results

3.1. Data

Our data set includes the Standard and Poor's S&P 500 stock market index (S&P500), and West Texas Intermediate (WTI) spot crude oil price. Data are monthly and cover the period September 1859 through December 2013. Data are extracted from the *Global Financial Database*. We seasonally adjust the both the data series using Census X13. Since we predict stock returns, we define stock returns (SR) as the month-on-month percentage change in the S&P 500 index, thus our effective sample starts from October 1859, which gives us a total of 1,851 observations. The starting and end points of our data are governed by data availability at the time of writing this paper. Note that for the sake of simplifying, we define the following variables, to which we refer to regularly in our discussion of the results: the natural logarithm of oil price (OP), oil price returns (OR) obtained as the month-on-month percentage change in the oil price, and the positive (OP^+) and negative (OP^-) components of the oil price, respectively. Our oil price data is unique in that it goes back in time, as far back as possible. The first oil well was drilled in the US on August 27, 1859 in Titusville, Pennsylvania. This defines the beginning of the modern era in the petroleum industry. We have data from this time period.

At this stage, it is important to emphasize that our decision to use nominal variables rather than real variables is motivated by two factors: First, consumer price index data at monthly frequency is only available from 1876 and, hence, by using real data we would have lost more than sixteen years of data, including the so-called first oil-shock over 1862-1864 due to the US civil war (Hamilton, 2011). Given our aim of understanding the historical relationship between US stock returns and oil prices, the longest possible sample period is needed and we do not compromise on this. Our second reason is motivated by the work of Hamilton (2011), who suggests that using real values induces measurement errors. More specifically, he writes: "... deflating by a particular number, such as the CPI, introduces a new source of measurement error,

which could lead to a deterioration in the forecasting performance. In any case, it is again quite possible that there are differences in the functional form of forecasts based on nominal instead of real prices” (page 370).

3.2. Preliminary Results

We first examine the basic statistics, namely, predictor persistency and endogeneity, and model heterokedasticity, on which the predictive regression model’s estimators are predicated. The results are reported in Table 1. Specifically, in columns 2 and 3, we report results from a test of the null hypothesis of a unit root in SR , OP , OP^+ and OP^- . For this purpose, we use the familiar augmented Dickey-Fuller (ADF, 1981) test, which, in the case of OP , OP^+ and OP^- , includes a time trend and an intercept, while in the case of SR , it includes only an intercept. The test statistic and the p -value are reported for each series, as is the estimated lag length, which is obtained by using the Schwarz Information Criteria (starting with a maximum of twenty-four lags). According to the ADF test, the unit root null is rejected for SR and OP^- , but not for OP and OP^+ . However, since the rejection of the null does not imply that OP^- is not persistent, in column 4 of Table 1, we also report the estimated first-order AR coefficient of OP , OP^+ and OP^- . As expected, for OP^- for which the null is rejected, and similar to OP and OP^+ for which the null is not rejected, the AR coefficient is very close to one. We, therefore, treat these variables as if they are unit root non-stationary. In the final three columns of Table 1, we report the mean, standard deviation, skewness, kurtosis and the Jarque-Bera (J-B) normality test, of SR , OP , OP^+ and OP^- . In summary, we find that while SR is more volatile than OP , OP^+ is more volatile than the other three series. Moreover, all four series are found to be non-normal.

[INSERT TABLE 1 ABOUT HERE]

We now turn to the results for heteroskedasticity reported in Table 2. With regard to autocorrelation, we report the Ljung-Box Q -statistic at the lag length of twelve for the squared values of SR , OP , OP^+ and OP^- . The results suggest that the null hypothesis of no autocorrelation must be rejected at the 1 percent level for all four series. While the evidence of autocorrelation in squared variables is indicative of ARCH, we also perform a formal Lagrange multiplier (LM) test for heteroskedasticity. Our approach here is as follows: We filter each of the four series through a twelfth-order autoregressive (AR) model. The LM test is then applied to the resulting filtered variables. According to the results (see Table 2), the null hypothesis of no ARCH is again strongly rejected for the four series under consideration.

[INSERT TABLE 2 ABOUT HERE]

As the final part of the preliminary analysis, we search for any evidence that oil price is endogenous. This is done and results are reported in Table 3; specifically, we report the OLS estimates of θ in Equation (3). We find, like the strong evidence that OP , OP^+ and OP^- are persistent and that SR and OP , OP^+ and OP^- are characterised by ARCH, that OP , and OP^- are strongly endogenous. The endogeneity of OP^+ , on the other hand, cannot be rejected at the 10 percent level of significance.

The main message emerging from the preliminary analyses of the data is that there is ample evidence that OP , OP^+ and OP^- are strongly persistent and endogenous, and the proposed predictive regression model is heteroskedastic. This implies the need to address these issues in estimating the predictive regression models, and it is these features that motivate us to use the

procedure of Westerlund and Narayan (2012), and to some extent also the procedure proposed by Lewellen (2004).

[INSERT TABLE 3 ABOUT HERE]

3.3. *Predictability Test Results*

We now turn to the results on in-sample predictability. In Table 4, we report results for forecast horizon (h); that is, for one-month ahead forecast ($h = 1$). We report the asymptotic GLS and Lewellen (2004) 95 percent confidence intervals for β . For the GLS test, we also report the sub-sample based confidence intervals. Looking first at the results from the Lewellen (2004) test, we find no evidence that stock returns are predictable when using OP , OP^+ and OP^- as predictors. One possible explanation for this weak evidence of predictability is the inability of the Lewellen (2004) OLS-based estimator to account for heteroskedasticity, which, as we documented earlier, is a strong feature of our data. Indeed, Westerlund and Narayan (2012) show that in the presence of heteroskedasticity the performance of the Lewellen (2004) test (in terms of its power to reject the null) weakens substantially. When we consider results from the GLS-based test, according to the asymptotic confidence intervals, there is evidence of stock return predictability resulting from the use of all three predictors; OP , OP^+ and OP^- . However, as mentioned earlier in Section 2.2, perhaps the most suitable confidence intervals with which to judge predictability are the sub-sample ones. The evidence of stock return predictability is robust to the application of sub-sample-based confidence intervals.²

[INSERT TABLE 4 ABOUT HERE]

²While we compute in-sample results for four horizons (one-month, 3-months, six-months and twelve-months), we only reported the one-month ahead results to save space. Further, our primary focus in this paper is the out-of-sample predictive ability of oil price and its component for stock returns, given the widely held view that predictive models require out-of-sample validation (Rapach and Zhou, 2013). Our in-sample results are, however, robust across the other horizons and are available upon request from the authors.

We now turn to the other main focus of the paper, i.e., the out-of-sample forecasting evaluation. Here, the performance of the oil price and its positive and negative components-based forecasting models are tested against the historical average-based forecasts, which is nothing but forecasts generated using Equation (1) with β set to zero. Following Narayan *et al.* (2014), we use half the sample (50%) to generate the first forecast. We believe that the usage of 50% of the observations (i.e., till 1936:11) allows us to cover in the out-of-sample all the major oil shocks in the history of the WTI oil price (Hamilton, 2011).³

The results are reported in Table 5. We begin by examining the relative Theil U statistic, which is the ratio of the Theil U from the unrestricted (oil price-based) model to the restricted (historical average) model. This implies that when the relative Theil U is less than one, the forecasts from the unrestricted model are better than those obtained from the restricted model. We report results for $h = 1, 3, 6$ and 12 . When using OP as a predictor, our unrestricted predictive regression model consistently outperforms the historical average model at all horizons, with the highest gain observed at a twelve-month-ahead forecast. Interestingly, the forecasting gain no longer exists with OP^+ except at the 12-month-ahead forecasting horizon, suggesting that positive oil shocks only play a role in predicting stock returns at a longer-horizon.

When using OP^- as a predictor, the results are reversed, with gains observed for shorter horizons ($h=1$ and 3) relative to the historical average. When we compare the relative Theil U , it has smaller values when the predictor is OP^- at shorter horizons compared to when the predictor is OP . However, with the OP predictor, we obtain higher gains at $h = 12$ compared to

³Our results are however, in general, robust to the choice of the size of the out-of-sample period, with greater predictability being observed for smaller out-of-sample horizons. Details of the results from using an out-of-sample period which involves 25 percent and 75 percent of the data are available upon request from the authors.

when OP^+ is the predictor. In sum, disaggregating the oil shock, especially into its negative component, seems to matter in predicting stock returns between one-month and a quarter year.

We consider two other measures of forecasting performance, namely, the out-of-sample R^2 (OOS_R) statistic proposed by Campbell and Thompson (2008), and the $MSE-F$ statistic suggested McCracken (2007). The OOS_R statistic is given by $OOS_R = 1 - (M\hat{S}E_1 / M\hat{S}E_0)$, where $M\hat{S}E_1$ and $M\hat{S}E_0$ are the mean square errors (MSE) of the out-of-sample predictions from the unrestricted and restricted models, respectively. Hence, $OOS_R > 0$ suggests that the unrestricted model outperforms the restricted model. The results reported in Table 5 mirror the findings obtained earlier based on the relative Theil U statistic, with the only difference being that it does not indicate any gain at $h=3$ from the predictive regression when OP^+ is the predictor compared to when OP is the predictor. Finally, the $MSE-F$ statistic tests the null hypothesis that the restricted and unrestricted models have equal forecasting ability. The null is tested against the one-sided alternative hypothesis that the MSE for the unrestricted model forecasts is less than the MSE for the restricted model forecasts. Formally, the statistic is given as:

$$MSE-F = (T-R-h+1) \cdot \bar{d} / M\hat{S}E_1 \quad (7)$$

where T is the total sample, R is number of observations used for estimation of the model from which the first forecast is formed (i.e. the in-sample portion of the total number of observations), $M\hat{S}E_i = (T - R - h + 1)^{-1} \sum_{t=R}^{T-h} (u_{i,t+1})^2$, $i = 1,0$, $\bar{d} = M\hat{S}E_0 - M\hat{S}E_1$, with u_i being the forecast error. A positive and significant $MSE - F$ statistic indicates that the unrestricted model forecasts are statistically superior to those of the restricted model. As can be seen from Table 5, for all cases where the relative Theil U is less than 1, i.e., $h = 1, 3, 6$ and 12 for OP ; $h = 12$ for OP^+ ; and $h = 1$ and 3 for OP^- , the $MSE-F$ statistic is positive and significant at

least at the 5 percent level of significance. Thus, overall, we find evidence of statistically significant forecasting gains in using oil prices as a predictor for S&P500-based US stock returns at all horizons relative to the historical average. In addition, disaggregating oil price shocks into its positive and negative components yields statistically significant forecasts relative to the benchmark at both long and short-horizons.

[INSERT TABLE 5 ABOUT HERE]

4. Concluding Remarks

In this paper, we re-visit the relationship between the oil prices and stock returns using an extensive and rich historical data that is monthly and begins in 1859; we, therefore, have over 150 years of monthly data. This type of historical analysis has not been previously undertaken. We find that oil price is a persistent and endogenous predictor variable and that our proposed stock return predictability model is heteroskedastic. Using a generalized least squares estimator of the null hypothesis of no predictability, we discover that the null hypothesis of no stock return predictability is rejected. We further disaggregate oil price into positive and negative oil prices. This allows us to test whether oil price nonlinearly predicts US stock returns. We find evidence of nonlinear predictability, that is, negative oil prices predict US stock returns more than do positive oil prices. We conclude by showing that US stock return predictability is robust on two fronts: first, in-sample evidence is corroborated by out-of-sample evidence; and, second, in out-of-sample evaluations our proposed oil price-based predictive regression model beats a constant returns model regardless of the choice of the in-sample period.

There are two key implications from our work. The first implication relates to the historical nature of the relationship between oil price and stock returns. While to-date using short

spans of data it was known that oil prices predicted stock returns, we showed that the same predictability relationship holds even when we consider this relationship over a sample period that spans 150 years. This means that the relationship between oil price and stock markets is not a new one. The second implication relates to the nonlinear nature of the relationship between oil price and stock price. We show that negative oil prices move stock returns more than do positive oil prices. This means that negative shocks emanating from the oil market are relatively more important for predicting stock returns. These implications should guide future empirical and theoretical work on these two markets.

References

- Asteriou, D., and Bashmakova, Y., (2013) Assessing the impact of oil returns on emerging stock markets: A panel data approach for ten Central and Eastern European countries, *Energy Economics*, 38, 204-211.
- Campbell, J.Y., (2008) Viewpoint: estimating the equity premium, *Canadian Journal of Economics*, 41, 1–21.
- Campbell, J.Y., and Thompson, S.B., (2008) Predicting excess stock returns out of sample: can anything beat the historical average? *Review of Financial Studies*, 21, 1509–1531.
- Casassus, J., and Higuera, F., (2012) Short-horizon return predictability and oil prices, *Quantitative Finance*, 12, 1909-1934.
- Cunado, J., and de Gracia, F.P., (2014) Oil price shocks and stock market returns: Evidence for some European countries, *Energy Economics*, 42, 365-377.
- Driesprong, G., B. Jacobsen, and Maat, B., (2008) Striking Oil: Another Puzzle? *Journal of Financial Economics*, 89, 307-327.
- Elyasiani, E., Mansur, I., and Odusami, B., (2011) Oil price shocks and industry stock returns, *Energy Economics*, 33, 966-974.
- Granger, C.W., and Yoon, G., (2002) Hidden cointegration. Economics Working Paper. University of California.
- Gupta, R., and Modise, M.P., (2013) Does the source of oil price shocks matter for South African stock returns? A structural VAR approach, *Energy Economics*, 40, 825-831.
- Hamilton, J.D., (2011) Nonlinearities and the macroeconomic effects of oil prices, *Macroeconomic Dynamics*, 15, 472–497.

Hong, H., W. Torous, and Valkanov, R., (2007) Do Industries Lead Stock Markets? *Journal of Financial Economics*, 83, 367-396.

Hong, H., and Stein, J.C., (1999) A unified theory of underreaction, momentum trading and overreaction in asset markets, *Journal of Finance*, 54, 2143-2184.

Lee, C-C., and Zeng, J-H., (2011) The impact of oil price shocks on stock market activities: Asymmetric effect with quantile regression, *Mathematics and Computers in Simulation*, 81, 1910-1920.

Lee, B-J., Yang, C.W., and Huang, B-N., (2012) Oil price movements and stock markets revisited: A case of sector stock price indexes in the G-7 countries, *Energy Economics*, 34, 1284-1300.

Lewellen, J., (2004) Predicting returns with financial ratios, *Journal of Financial Economics*, 74, 209–235.

McCracken, M.W., (2007) Asymptotics for out-of-sample tests of Granger causality, *Journal of Econometrics*, 140, 719–752.

Narayan, P. K., and Sharma, S.S., (2011) New Evidence on Oil Price and Firm Returns, *Journal of Banking and Finance*, 5, 3253-3262.

Rapach, D. E., and Zhou, G., (2013) Forecasting stock returns, in Graham Elliott and Allan Timmermann, eds: *Handbook of Economic Forecasting, Volume 2* (Elsevier, Amsterdam).

Westerlund, J., and Narayan, P.K., (2012) Does the choice of estimator matter when forecasting returns? *Journal of Banking and Finance*, 36, 2632–2640.

Westerlund, J., and Narayan, P.K., (2014) Testing for Predictability in Conditionally Heteroskedastic Stock Returns. *Journal of Financial Econometrics*, <http://dx.doi.org/10.1093/jfinec/nbu003>.

Table 1: Selected descriptive statistics of the data

var.	Unit root test		AR(1)	Mean	Std.Dev	Skew.	Kurt.	JB
	Test stat	Lags						
<i>SR</i>	-9.52***	24		0.38	0.38	-0.53	11.59	5780.1***
<i>OP</i>	-2.55	18	0.9988***	1.40	1.32	0.72	2.79	161.5***
<i>OP⁺</i>	-5.39***	17	0.9990***	28.00	10.07	-0.60	3.08	111.5***
<i>OP⁻</i>	-3.68***	24	0.9981***	-28.68	9.24	0.92	3.56	283.1***

Notes: See note to Figure1; *** p -value is <0.01 ; Lag-lengths chosen based on *SIC*; Std. Dev: Standard Deviation; Skew.:Skewness; Kurt.: Kurtosis; J-B test is the Jarque-Bera test of normality.

Table 2: Results for Ljung-Box Q -statistic and ARCH F -statistics

Variables	ARCH F -stat	Variables	Ljung-Box Q -stat
SR	4.6864***	$(SR)^2$	1103.8***
OP	4.6212***	$(OP)^2$	20602***
OP^+	2.5080***	$(OP^+)^2$	21589***
OP^-	9.1263***	$(OP^-)^2$	21604***

Note: See note to Figure 1; *** p -value is < 0.01 .

Table 3: Results for endogeneity test

Variables	Coefficient	<i>t</i> -stat	<i>p</i> -value
<i>OP</i>	3.0568	2.4909	0.0128
<i>OP</i> ⁺	3.2549	1.6772	0.0937
<i>OP</i> ⁻	4.2389	2.1949	0.0283

Note: See note to Figure 1.

Table 4: 95% confidence intervals for β , $h=1$

Variables	$t_{\text{sub_FGLS}}$	t_{FGLS}	t_{LEW}
OP	[0.028, 0.016]	[0.012, 0.012]	[-0.0012, 0.027]
OP^+	[0.127, 0.103]	[0.066, 0.066]	[-0.0062, 0.104]
OP^-	[0.045, 0.076]	[0.017, 0.018]	[-0.081, 0.054]

Note: See note to Figure 1; t_{FGLS} , t_{LEW} , $t_{\text{sub_FGLS}}$ corresponds to 95% confidence intervals for β in Equation (4), corresponding to the Westerlund and Narayan (2012) GLS estimation, Lewellen (2004) OLS estimation, and the sub-sample based GLS estimation of Westerlund and Narayan (2012), respectively; h stands for forecast horizon.

Table 5: Out-of-sample forecast evaluation

Variables	$h=1$			$h=3$			$h=6$			$h=12$		
	Rel. Theil U	OOS_R	$MSE-F$	Rel. Theil U	OOS_R	$MSE-F$	Rel. Theil U	OOS_R	$MSE-F$	Rel. Theil U	OOS_R	$MSE-F$
OP	0.994	0.013	11.77***	0.994	0.013	11.76***	0.994	0.013	11.55***	0.989	0.023	20.92***
OP^+	1.028	-0.057	-52.42	1.019	-0.038	-35.25	1.020	-0.041	-37.62	0.999	0.002	2.01*
OP^-	0.992	0.016	14.82***	0.993	0.013	12.06***	1.000	-0.001	-10.42	1.036	-0.073	-66.99

Note: See note to Figure1; Rel. Theil U : Relative Theil U ; OOS_R : out-of-sample R^2 statistic of Campbell and Thompson (2008); $MSE-F$: $MSE-F$ statistic of McCracken (2007); *** p -value is <0.01 ; * p -value is <0.10 .