Energy Economics 41 (2014) 137–146

Contents lists available at ScienceDirect

Energy Economics

journal homepage: www.elsevier.com/locate/eneco

Do oil prices predict economic growth? New global evidence

Paresh Kumar Narayan [a,*], Susan Sharma [a], Wai Ching Poon [b], Joakim Westerlund [a]

[a] Centre for Financial Econometrics, School of Accounting, Economics and Finance, Deakin University, Australia
[b] School of Business, Monash University, Malaysia Campus, Malaysia

ARTICLE INFO

Article history:
Received 11 June 2013
Received in revised form 7 October 2013
Accepted 7 November 2013
Available online 21 November 2013

JEL classification:
C22
E31
E37
F43

Keywords:
Economic growth
Predictability
Oil price

ABSTRACT

In this paper, we test whether oil price predicts economic growth for 28 developed and 17 developing countries. We use predictability tests that account for the key features of the data, namely, persistency, endogeneity, and heteroskedasticity. Our analysis considers a large number of countries, shows evidence of more out-of-sample predictability with nominal than real oil prices, finds in-sample predictability to be independent of the use of nominal and real prices, and reveals greater evidence of predictability for developed countries.

1. Introduction

That there is a relationship between oil price and economic growth is well-known. Two strands of the literature have reinforced this. Consider first the studies that have estimated the effects of oil prices on economic growth.[1] The main findings of this literature are two-fold: (a) oil price generally has a negative effect on economic growth (Kilian, 2008; Kilian and Vigfusson, 2011a); and (b) the oil price effect need not be linear (Hamilton, 2003; Kilian and Vigfusson, 2011b). The latter finding implies that oil prices tend to affect countries differently depending on their stage of development. The second strand of literature owes much to the early work of Hamilton (1983), and tests whether oil prices have any predictive content. Typically, these studies fit a predictive regression model of economic growth in which oil price appears as a predictor variable; see also Hamilton (2011).

As much as this literature is growing and is attractive, given the gradual rise in oil prices over the last decade and the ramifications for economic performance, a key limitation is also rather obvious. Much of the research on the economic growth–oil price nexus focuses on the US economy. Outside of the US, not much is known on whether or not the oil price predicts economic growth. In light of this research gap, we test whether oil price predicts economic growth in 45 countries, of which 28 are developed and 17 are developing. We use quarterly time series data. Our predictive regression model is familiar in that economic growth (proxied by either growth in real gross domestic product or industrial production) is regressed on the one-period lagged oil price variable.

The contribution of our paper is three-fold. First, our paper not only focuses on the US, which has previously been the main subject of this literature, but also includes as many as 44 additional developed and developing countries. A multi-country study of whether or not oil price predicts economic growth allows us to better understand the role of oil prices on a more global level. At this stage, it is fair to claim that the role of oil prices in economic growth is very much unknown from a global point of view. Our proposed empirical investigation narrows this research gap.

Our second contribution is relatively more methodological in that we pay particular attention to the salient features of data, namely, persistency, endogeneity, and heteroskedasticity, that matter directly for the performance of predictive regression models. The first issue relates to the persistent nature of the predictor variable. Specifically, the existence of persistent predictors has been shown to lead to the failure of conventional asymptotic theory for exogenous regressors (see Elliot and Stock, 1994), leading to deceptive inference. We find that oil price is highly persistent; not only do we accept the unit root null hypothesis, we also find the autoregressive coefficient of the oil price variable to be close to one. Then there is the issue of endogeneity of the predictor

* Corresponding author at: Alfred Deakin Professor, School of Accounting, Economics and Finance, Faculty of Business and Law, Deakin University, 221 Burwood Highway, Burwood, Victoria 3125, Australia. Tel.: +61 3 9244 6180; fax: +61 3 9244 6034.
E-mail address: paresh.narayan@deakin.edu.au (P.K. Narayan).

[1] Indeed there are studies, such as Ferderer (1996), Elder and Serletis (2010), Rahman and Serletis (2012), that have considered the effects of oil price volatility on economic growth.

0140-9883/$ – see front matter
http://dx.doi.org/10.1016/j.eneco.2013.11.003

variable. It would be bold to claim that oil price is purely exogenous. For example, by generating higher demand for oil, growth could also influence oil price. Therefore, formally testing whether or not oil price is endogenous is a matter of prerequisite, for, as already alluded, endogeneity of the predictor variable has been shown to bias the results on predictability. The final issue is that of heteroskedasticity. Westerlund and Narayan (2012) show that if heteroskedasticity is present and it is correctly accounted for in predictive regression models, then the properties of the resulting predictability test are better compared to when heteroskedasticity is ignored. Our approach to addressing these three issues is to use the bias-adjusted ordinary least squares (OLS) estimator of Lewellen (2004), and the Westerlund and Narayan (2012) generalised least squares (GLS) estimator. The main difference between the two is that while the former estimator accounts for only persistency and endogeneity, the latter estimator is flexible enough to cater for all three features of the data. To put these issues into perspective, let us at the outset acknowledge that while the literature has been mindful of the issue of persistent predictors, the issue of endogeneity has received little attention, while that of heteroskedasticity has been completely ignored. Ignoring these features of the data comes at a cost as they have direct implications for the outcome on predictability.

Third, we establish the robustness of our findings by undertaking both in-sample and out-of-sample predictability analyses. This approach is not common in the literature, and some studies (Ashley et al., 1980; Rapach and Wohar, 2006) suggest that perhaps an out-of-sample analysis is relatively more important to policy makers than in-sample evidence. A related group of studies (Foster et al., 1997; Lo and MacKinlay, 1990) claim that in-sample tests suffer from data mining. Inoue and Kilian (2004), however, show that in-sample and out-of-sample tests of predictability are equally reliable against data mining under the null hypothesis of no predictability. What is clear from this literature is that there is no shortage of tension when it comes to the choice between in-sample and out-of-sample evaluations and we are avoiding being caught in this debate. The best way forward is to undertake both in-sample and out-of-sample evaluations. Doing so not only makes the predictability analysis complete but it also allows us to gauge the robustness of our results.

Briefly foreshadowing the main findings, we find that nominal oil price predicts economic growth for 37 of the 45 countries and for around 70% of the countries there is evidence of out-of-sample predictability. When we use real oil price, like with nominal oil price, we discover strong evidence of in-sample predictability (for 36 countries). However, evidence on out-of-sample predictability is weak. At best, only for around 55% of the countries there is evidence of out-of-sample predictability. Finally, we find that with nominal oil price both in-sample and out-of-sample evidence of predictability are found for 33 countries while for real oil price this evidence is only found for 30 countries.

We organise the balance of the paper as follows. In Section 2, we discuss the data and methodology. In Section 3, we discuss the results. In the final section, we provide concluding remarks.

2. Data and methodology

2.1. Data

This paper is based on a quarterly data set that includes 45 countries. Of these 45 countries, 17 are developing countries and the balance is developed countries. The sample size is dictated by data availability. We have quarterly data. For 68% of the countries in our sample the data span the period 1983Q2 to at least 2010Q4. Therefore, for most countries we have no less than 113 quarterly observations. The specific dates of data for each country are reported in the last column of Table 1. The world average crude oil price and industrial production index are obtained from the *International Financial Statistics* (IFS) published by the International Monetary Fund, while data on quarterly real GDP growth rate are obtained from the *World Development Indicators* published by the World Bank. The nominal crude oil price was converted into the real crude oil price by using the country-specific consumer price index, which was obtained from the IFS.

2.2. Estimation approach

A typical predictive regression model, where oil price is considered as a predictor of economic growth, has the following form:

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

Here, y_t is the economic growth in quarter t proxied by either the growth rate in real GDP or industrial production, and OP_t is the average world crude oil price in US dollars in the same quarter. 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. If it is, one can expect a bias, leading to deceptive inference on the no predictability null. Given that in our empirical analysis we have relatively small sample sizes, the implications of endogeneity could be serious. To avoid this, we follow Westerlund and Narayan (2012) and model oil price as follows:

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

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

$$\varepsilon_{y,t} = \theta\varepsilon_{op,t} + \epsilon_t \tag{3}$$

where ϵ_t is again mean zero and with variance σ^2_ϵ.

We use two estimators, bias-adjusted OLS and GLS. Both estimators are based on making Eq. (1) conditional on Eq. (2), thereby removing the effect of the endogeneity. The resulting conditional predictive regression can be written as[2]:

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

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 Eq. (4).

The key difference between this estimator and the one of Westerlund and Narayan (2012) 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.[3] In particular, it is assumed

[2] The literature on predictability models has moved away from treating a predictor variable as purely stationary because in practice it is not. Although the null hypothesis of unit root can be comfortably rejected for many predictors, they are still very highly persistent. In other words, many predictors are shown and, as a result, known to be only slowly mean-reverting. Let us see this. Denote the predictor variable by y_t, such that we have $y_t = \rho y_{t-1} + \varepsilon_t$. Standard asymptotic theory, which presumes that $|\rho| < 1$, is likely to be inappropriate because predictors are shown to be persistent even though the unit root null hypothesis can be comfortably rejected (see Campbell and Yogo, 2006; Elliot and Stock, 1994; Lewellen, 2004; Westerlund and Narayan, 2012). In particular, Elliot and Stock (1994) show that even if $\{t_t\}_{t=1}^T$ is stationary, if $\rho \approx 1$, the standard asymptotic theory is likely to provide a poor approximation in small samples. As a response to this, researchers have considered alternative frameworks based on 'local asymptotic theory' (see Campbell and Yogo, 2006; Cavanagh et al., 1995; Lanne, 2002; Lewellen, 2004; Torous et al., 2004; and Westerlund and Narayan, 2012). This theory allows one to model a highly persistent predictor variable.

[3] The size adjusted power gain from using the GLS test statistic over the OLS test statistic in small sample sizes is estimated to be around 20% (Westerlund and Narayan, 2012).