The Quarterly Review of Economics and Finance 71 (2019) 14–26

Contents lists available at ScienceDirect

The Quarterly Review of Economics and Finance

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

Institutions and return predictability in oil-exporting countries[☆]

Sirio Aramonte[a,*], Mohammad R. Jahan-Parvar[b], Justin K. Shugarman[c]

[a] The Bank for International Settlements, Centralbahnplatz 2, Basel 4002, Switzerland
[b] The Board of Governors of the Federal Reserve System, 20th St. and Constitution Ave. NW, Washington, DC 20551, USA
[c] New York University, 19 W. 4th Street, New York, NY 10012, USA

ARTICLE INFO

Article history:
Received 27 September 2017
Received in revised form 31 July 2018
Accepted 18 September 2018
Available online 28 September 2018

JEL classification:
G12
O16
O50

Keywords:
Country studies
Delayed reaction
Equity market return predictability
Institutional quality

ABSTRACT

We study whether oil price changes have predictive power for stock market returns in oil-exporting countries, and we investigate the link between predictability and the quality of each country's institutions. Returns are predictable for half the countries we consider, and predictability is stronger when institutional quality is lower. We argue that the relation between predictability and institutional quality reflects the preference of countries with weaker institutions to consume oil windfalls locally rather than smooth out their impact by, for example, investing the proceeds internationally.

1. Introduction

In this paper, we investigate whether oil price changes can predict short-term stock returns in oil-exporting countries. This predictability stems from delayed reaction to changes in oil price information and we link the strength of the predictive relation to institutional quality, for which we use Transparency International's Corruption Perceptions Index as a proxy. We find that market index returns are predictable in half the countries we consider, and that predictability is stronger for countries with weaker institutions. Thus, we join evidence from two strands of the literature in our study, one from asset pricing theory and another from political economy.

Our analysis builds on two key assumptions. First, that due to the delayed response of stock returns to oil price changes (identified by Driesprong, Jacobsen, & Maat, 2008), stock returns in t will respond to oil price changes in $t-1$. This analysis builds on key insights introduced by Hong and Stein (1999). They develop a theoretical framework for slow diffusion of information in financial markets, due to bounded rationality of agents who do not react to varied sources of information in real time, leading to delayed reaction in asset prices in one market (say, equities) to information stemming from other markets such as oil in our study. Hong, Torous, & Valkanov, 2007, Driesprong et al. (2008), and Fan and Jahan-Parvar (2012) investigate these insights empirically.[1]

Second, based on the extant literature, we posit that institutional quality is negatively related to the propensity to consume oil windfalls, directly through pro-cyclical fiscal policies or indirectly through inefficient expenditure. In addition, we posit that return predictability reflects the extent to which oil windfalls are

☆ We thank Hadi Salehi-Esfahani (the Editor), two anonymous referees, David Arseneau, Cecilia Caglio, Luca Guerrieri, Lutz Kilian, M. Hashem Pesaran, and seminar participants at the 2015 Georgetown Center for Economic Research Biennial Conference and at the Federal Reserve Board for helpful comments. All remaining errors are ours. This article represents the views of the authors, and should not be interpreted as reflecting the views of the Bank for International Settlements or of the Board of Governors of the Federal Reserve System, or other members of their staff.

* Corresponding author.
E-mail address: sirio.aramonte@bis.org (S. Aramonte).

[1] The $t-1$ lag refers to our use of equity returns in a given month and oil price changes in the previous month. However, we build these monthly values from daily data: the slow diffusion of information theory implies that information stemming in another market/sector should get impounded in prices of the asset under evaluation within a few days. We test this by skipping an increasing number of trading days (between the end of month $t-1$ and the beginning of month t) when we construct the oil price change measures.

https://doi.org/10.1016/j.qref.2018.09.002
1062-9769/

consumed locally rather than, for instance, invested overseas to attenuate the impact of the windfall on the local economy that the strength of the predictive regression depends on institutional quality, because lower institutional quality means that a country is less likely to smooth the impact of changes in oil prices on government revenue. Analyzing the interaction between economic growth and institutional quality is typically complicated by endogeneity issues. Our analysis assumes that the institutional quality impacts fiscal policy, and not vice versa. We do so since fiscal policy is directly controlled by policymakers, which reduces reverse-causality concerns. Our second assumption is based on the findings of a large literature. Persson (2002) highlights that political institutions influence economic choices, including "fiscal policy, broadly defined to include rents sought by corrupt politicians" (page 884). Fatás and Mihov (2003) show that countries that deploy more prudent fiscal policies face lower output volatility and higher economic growth, and that such countries are characterized by majoritarian and constrained political systems. Alesina, Tabellini, and Campante (2008) find evidence that corrupt democracies have more pro-cyclical fiscal policies. Pieschacon (2012) analyzes the role of fiscal policy on how oil shocks affect consumption and relative prices in oil-exporting small open economies through a dynamic stochastic general equilibrium (DSGE) model. She shows that countries like Norway, which invests oil windfalls internationally, see smaller fluctuations in domestic variables like output or the price of non-tradeable goods in response to an increase in oil revenues than countries like Mexico, whose fiscal policies allow consumption and other domestic variables to fluctuate more closely with oil revenues.

The existing literature on whether oil prices can predict stock returns is largely focused on developed markets. Huang, Masulis, and Stoll (1996) study the lead-lag relation between U.S. stock returns and returns on oil futures. With the exception of oil companies, they find weak correlations. In an analysis of how oil prices interact with changes in expected returns and cash flows, Jones and Kaul (1996) show that lagged quarterly changes in oil prices can, in isolation, predict quarterly stock-market returns for the United States, Canada, and Japan. Reboredo and Rivera-Castro (2014) study the relation between daily stock returns and oil prices in Europe and the United States, and find no evidence of a relation before the 2008 financial crisis. Driesprong et al. (2008) conduct an extensive predictability study and highlight that oil prices have weak predictive power in the case of developing economies. Kilian and Park (2009) show that the implications of oil prices shocks for U.S. stock returns depend on whether prices have changed in response to supply or demand shocks, with the latter explaining a larger fraction of stock returns. Gogineni (2010) observes that the correlation between industry-level stock returns and oil price changes varies according to cost-side and demand-side dependence to oil of the firms in each industry. In this paper, we follow Jones and Kaul (1996) and Driesprong et al. (2008), and remain agnostic as to whether the predictability is driven by demand or supply shocks.

In addition to bridging the literature on return predictability and the interaction of fiscal policy and institutional quality, our study is indirectly related to the literature on the predictability of the business cycle using oil prices, and to studies of the "resource curse," which focus on how the economic development of resource-rich countries interacts with weak institutions. Studies of return predictability focus on the short term, but oil prices also have implications at the business-cycle frequency on macroeconomic variables such as output or unemployment. Hamilton (1983) and Hamilton (1996) highlight that shocks to oil prices often precede recessions in the United States. Barsky and Kilian (2004) emphasize that oil prices are endogenous with respect to economic conditions in the United States, and that, while oil shocks likely contribute to macroeconomic fluctuations, their role is not necessarily pivotal. See Hamilton (2011) and Kilian (2008) for recent literature overviews.

The long-run effect of oil income on the economic growth of oil-producing countries depends on a number of factors, chiefly the level and persistence of oil revenues (Esfahani, Mohaddes, & Pesaran, 2014). However, the institutional characteristics of a country also play a role, since they influence how efficiently these revenues are spent. Sachs and Werner (1995) find that the abundance of natural resources is negatively related to economic growth. Mehlum, Moene, and Torvik (2006) document that this "resource curse" is only found in countries characterized by inefficient institutions. Mauro (1995) also highlights that corruption and bureaucratic inefficiency significantly reduce both investment and economic growth. Leite and Weidmann (1999) and Vicente (2010) suggest that the causality can run in either direction, and that the availability of natural resources can reduce the efficiency of a country's institutions. James (2015) interprets the resource curse in terms of commodity price trends, since resource-rich countries grow more slowly when the commodity they depend on experiences price declines. See van der Ploeg (2011) for an encompassing review of the literature on this topic.

While the literature on the impact of corruption/institutional quality on economic growth generally acknowledges the endogeneity of the two variables of interest (e.g., Mauro, 1995), our analysis assumes that the institutional quality impacts fiscal policy, and not vice versa. The reason is that fiscal policy is a variable over which policymakers have direct control, which reduces reverse-causality concerns. Our empirical strategy is broadly comparable to Alesina et al. (2008), who study the pro-cyclicality of fiscal policy conditional on corruption, and use panel regressions in which corruption is not instrumented.

The rest of the paper is organized as follows: Section 2 describes the data and the empirical implementation; Section 3 presents the results; and Section 4 concludes.

2. Data and empirical implementation

2.1. Data

Our sample includes countries for which oil exports are a significant source of revenue, measured with the value of oil exports over gross domestic product (GDP). The value of oil exports is calculated by multiplying the average yearly price for a barrel of West Texas Intermediate oil (WTI, from the Federal Reserve Economic Data website) by the number of oil barrels exported in a given year (crude oil including lease condensate, from the website of the U.S. Energy Information Administration). Crude oil production figures are also from the U.S. Energy Information Administration. GDP in current U.S. dollars (USD) is from the World Bank (series NY.GDP.MKTP.CD).

Table 1 shows the countries for which (1) the oil revenue/GDP ratio is above 1% in 1995, 2003, and 2010, and (2) data on the market capitalization of listed companies to GDP (series CM.MKT.LCAP.GD.ZS from the World Bank) is available in the same years. Requiring that the ratio is above 1% in each of the three years ensures that oil revenues are a meaningful fraction of GDP throughout the sample. We consider the three years above because, as detailed later in this Section, the return series are mostly available from the mid-1990s. Finally, we only consider countries for which data on stock market capitalization to GDP is available to ensure that we can evaluate the role of heterogeneity in financial development. As shown in Section 3, our results are not driven by heterogeneity in financial market development.

Due to constraints on the availability of returns data, we exclude two of the countries in our sample (Iran and Trinidad and Tobago).

Table 1
Value of oil exports and of listed companies relative to gross domestic product (GDP)

Country	Oil exports to GDP (%)			Stock market to GDP (%)		
	1995	2003	2010	1995	2003	2010
Nigeria	40	36	18	7	14	14
Oman	38	40	35	14	23	34
Saudi Arabia	30	36	38	29	73	67
Kuwait	29	30	34	53	124	100
Iran	*19*	*19*	*16*	*7*	*25*	*20*
Venezuela	*16*	*21*	*12*	*5*	*5*	*1*
Norway	11	14	11	30	42	60
Trinidad and Tobago	*7*	*9*	*11*	*21*	*94*	*59*
Ecuador	7	9	15	11	7	8
Egypt	4	1	1	13	33	38
Russia	4	12	9	4	54	66
Malaysia	3	4	3	251	153	166
Indonesia	3	3	1	33	23	51
Tunisia	3	2	5	22	9	24
Mexico	3	3	4	26	17	43
Colombia	2	3	5	19	15	73
Canada	1	2	3	61	101	134

The table shows the value of oil exports and of listed companies relative to GDP, in percent. Annual exports of crude oil (including lease condensate) are from the U.S. Energy Information Administration, and they are multiplied by the average annual WTI price. Current GDP in U.S. dollars is from the World Bank (series NY.GDP.MKTP.CD). The ratio of the market capitalization of listed companies to GDP is provided by the World Bank (series CM.MKT.LCAP.GD.ZS). Countries are sorted on the basis of the value of oil exports to GDP in 1995. Countries for which we are unable to source stock return data are shown in italics.

In addition, we exclude Venezuela since its stock market is not well developed, as highlighted by the very low ratio of market capitalization to GDP. The underdeveloped nature of the Venezuelan market renders prices illiquid and unreliable. Hence, we study 14 oil exporting countries that cover the Middle East, Africa, the Americas, Asia, and Europe. The importance of oil exports relative to GDP varies significantly in the cross-section, ranging from 40% to the low single digits. Over time, the ratio is generally stable, dropping sharply only for Nigeria. There is significant heterogeneity in financial development, with Canada, Kuwait, and Malaysia having stock market capitalization in excess of GDP, and Ecuador and Nigeria having relatively small stock markets.

Our proxy for institutional quality is the Corruption Perceptions Index maintained by Transparency International. Since the acronym CPI is normally associated with the Consumer Price Index, in the remainder of the paper we refer to Transparency International's index as IQI, shorthand for Institutional Quality Index. The index ranks countries on the basis of the perceived corruption of their public sectors as it emerges from surveys of professional analysts, and it consolidates rankings from various independent institutions that analyze governance and business climate.[2] Wilhelm (2002) documents that IQI is highly correlated with burdensome levels or regulation and black market activity, two other measures of corruption in an economy. IQI data for all of the countries in our sample is available starting in 2003, and we measure the institutional quality of each country with the average IQI between 2003 and 2013. We should note that, while interpreting corruption measures as proxies for "institutional quality" is to some extent arbitrary, previous studies have already highlighted this connection in light of conceptual overlaps and the empirical challenges in distinguishing the two.[3]

We report the average IQI in the first column of Table 2. The index ranges between 0 and 100, with the latter indicating the lowest possible level of perceived corruption. There are large differences across countries, with Canada and Norway scoring high, and Nigeria and Russia at the other end of the spectrum. Interestingly, a comparison of Tables 1 and 2 shows that higher perceived corruption is not necessarily associated with higher dependence on oil exports. For instance, Oman is highly dependent on oil and has the third highest IQI score in the sample. In Table 3 we show that the corruption index is generally stable over time, which implies that we can use the average IQI as a proxy for institutional quality throughout the sample, following the approach of Alesina et al. (2008).

We obtain stock returns from Thomson Reuters' Datastream. Following Driesprong et al. (2008), we use MSCI indexes when available (see Table 2 for details and summary statistics). While our analysis is based on total return indexes, in which dividends are reinvested, excluding dividends has little effect on the results. All indexes are expressed in USD, and data is available through December 2017. Some of the series go back to at least the 1980s, however a fair number of them start in the mid-to-late 1990s. We ensure that the regression results are comparable across countries by only considering data from January 1995 if the series are available from an earlier date. Table 2 shows the starting dates for the various return series.

2.2. Empirical implementation

Our research design draws from the results in Driesprong et al. (2008), who conduct an extensive predictability study, and find that oil price changes can predict local stock returns for a large cross-section of countries. They conclude that the predictability is not due to a potential co-movement of oil prices with risk premia, but to delays in incorporating oil-related information into stock prices.

The key assumption of this paper is that, because of the delayed response identified by Driesprong et al. (2008), stock returns in t will respond to oil price changes in $t-1$. In addition, we posit that

[2] For additional details, see `https://www.transparency.org/research/cpi/overview`

[3] Mehlum, Moene, and Torvik (2006, pg.3) note that institutions conducive to rent-seeking behavior are characterized by, among other items, "weak rule of law, malfunctioning bureaucracy, and corruption". Conversely, the survey-based measure of corruption built by Vicente (2010) probes the quality of a wide range of public services, including courts of law and the bureaucracy, and Mauro (1995) suggests that evaluating the efficiency of the judiciary and bureaucratic systems can improve the measurement of corruption. Several measures of corruption are available. Mauro (1995) studies Business International indexes, while Alesina et al. (2008) and Fisman and Miguel (2007) use the corruption measure of Kaufmann et al. (2005) and Kaufmann et al. (2006). Fisman and Miguel (2007) find that this measure is highly correlated with the rankings of Transparency International that we use in this paper.

Table 2
Summary statistics

	IQI	Return statistics				Details on return time series		
		Mean	Std. Dev.	Skewn.	Kurtosis	Starting date	Provider	Ticker
Nigeria	22	0.67	8.81	−0.51	7.72	1995/07	S&P	IFGDNG$
Oman	54	0.64	5.41	−1.25	11.05	2000/05	S&P	IFGDOM$
Saudi Arabia	40	0.61	7.26	−0.76	4.85	1998/01	S&P	IFGDSB$
Kuwait	45	0.22	6.24	−0.26	5.10	2005/01	S&P	IFGDKW$
Norway	87	0.45	7.69	−1.26	8.11	1995/01	MSCI	MSNWAY$
Ecuador	25	−0.25	9.35	−1.82	22.45	1996/01	S&P	IFFMEC$
Egypt	31	0.76	9.43	−0.27	5.50	1995/01	MSCI	MSEGYT$
Russia	25	0.65	14.77	−1.08	9.65	1995/01	MSCI	MSRUSS$
Malaysia	49	0.08	7.87	−0.27	8.86	1995/01	MSCI	MSMALF$
Indonesia	26	0.27	12.31	−0.68	6.67	1995/01	MSCI	MSINDF$
Tunisia	44	−0.01	4.81	−0.27	5.82	1996/01	S&P	IFFMTU$
Mexico	34	0.52	8.01	−1.25	7.28	1995/01	MSCI	MSMEXF$
Colombia	37	0.63	9.02	−0.36	3.89	1995/01	MSCI	MSCOLM$
Canada	86	0.59	5.82	−1.09	7.23	1995/01	MSCI	MSCNDA$
S&P 500		0.38	4.50	−0.88	4.54			
WTI		0.75	9.73	−0.64	4.35			
Brent		0.84	10.66	−0.64	4.91			

The first column shows the average between 2003 and 2013 of the Institutional Quality Index (IQI). 2003 is the first year in which all the countries listed in the table are ranked by Transparency International. For the 15 countries we study, the remaining columns report summary statistics and other information for the monthly excess log-returns on stock indexes with reinvested dividends. The last three lines show statistics for monthly excess log-returns on the S&P 500 index, and monthly log-changes in the price of West Texas Intermediate (WTI) oil and Brent oil. Returns in percent. The sample ends in 2017.

Table 3
The Corruption Perceptions Index over time

	IQI		
Country	2003	2008	2013
Nigeria	14	27	25
Oman	63	55	47
Saudi Arabia	45	35	46
Kuwait	53	43	43
Norway	88	79	86
Ecuador	22	20	35
Egypt	33	28	32
Russia	27	21	28
Malaysia	52	51	50
Indonesia	19	26	32
Tunisia	49	44	41
Mexico	36	36	34
Colombia	37	38	36
Canada	87	87	81

The table shows IQI in three years for each country. Note that the index was reported on a scale from 1 to 10 before 2012, and the values in 2003 and 2008 have been multiplied by 10.

the strength of the predictive regression will vary on the basis of institutional quality, since lower institutional quality implies that a country is more likely to consume oil windfalls rather than smooth their business-cycle effects by, for instance, investing the proceeds through a sovereign wealth fund. Our assumption that institutional quality affects fiscal policy is based on the results in Alesina et al. (2008), who find that corrupt democracies have more pro-cyclical fiscal policies, and in Fatás and Mihov (2003), who show that constrained political systems lead to less volatile fiscal policies.

The main advantage of our approach is that we do not need to model cyclically adjusted fiscal policy, which, as noted by Fatás and Mihov (2003), is complicated by the simultaneity of output and fiscal policy. Instead, we posit a link between fiscal policy and equity returns' reaction to oil price changes, since the more changes in oil revenue affect public spending, the more local stock returns are affected by oil price changes. While our methodology is simple to implement, we need to rule out alternative explanations for the relation between institutions and predictability, which we do in the next section.

Fig. 1 shows a clearly positive relation between institutional quality and the size of each country's sovereign wealth fund.[4] The regression slope reported in the figure, which we can interpret as the elasticity of sovereign wealth funds' size relative to institutional quality, is equal to 3.96. While the low number of observations makes statistical inference problematic, the positive relation shown in the figure supports our assumption that higher institutional quality is associated with higher smoothing of oil prices's impact on the business cycle.

We emphasize that Fig. 1 is based on data that are qualitatively different from those we use when studying return predictability. In particular, sovereign wealth funds invest in foreign assets (Kotter and Lel, 2011), while we focus on local equity market returns. Additionally, sovereign wealth funds' assets are stock data, while equity market returns reflect investment flows in and out of stock markets. Finally, sovereign wealth funds' assets reflect past government decisions, while investment flows into the stock market reflect expectations of future economic conditions, including anticipated fiscal policies. Given that institutional characteristics are persistent, as shown in Table 3, they can explain the way the IQI variable is related to both return predictability and sovereign wealth funds' assets.

We need to clarify why we study the effect of oil price changes on future returns rather than contemporaneous returns. The reason is that oil price changes are correlated with stock market volatility, especially when the price of oil drops.[5] As a result, the effect of oil price changes on contemporaneous returns can be obfuscated by simultaneous changes in expected returns driven by fluctuation in aggregate risk, as reflected in higher volatility. Our research design allows us to obtain a cleaner estimate of the impact of oil price changes on stock prices, since aggregate equity market indexes incorporate general news about the state of the economy more

[4] Data on sovereign wealth funds' assets is from the Sovereign Wealth Fund Institute, as of January 2015. Fig. 1 does not show four of the fifteen countries we study, because data on sovereign funds' assets are unavailable. These countries are Colombia, Ecuador, Egypt, and Tunisia.

[5] Over 1995–2013, the correlation between monthly oil price changes and monthly S&P 500 return volatility is −28% (−40% if we only consider negative returns). Excluding October 2008, when both oil prices and stock prices had outsized changes, the correlations are −15% and −28%, respectively.

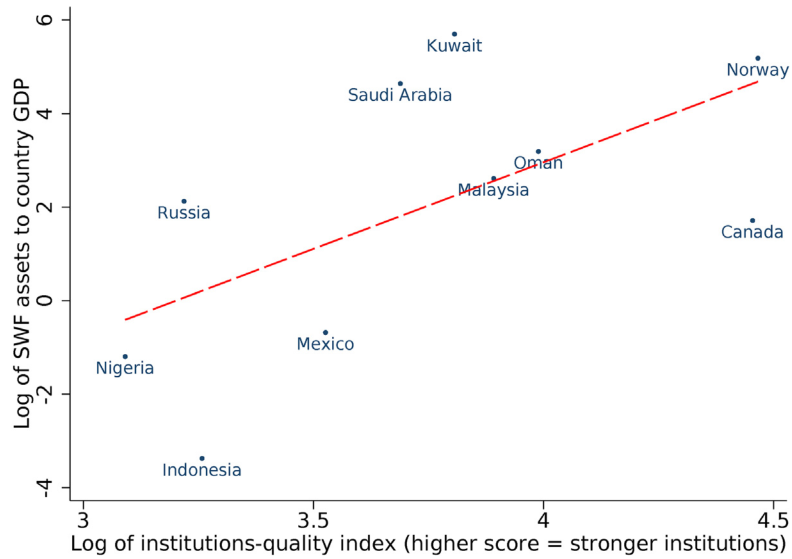

Fig. 1. Institution quality and the size of sovereign wealth funds. The vertical axis shows the logarithm of the percentage ratio of total assets held in sovereign wealth funds for a given country over the GDP of the country. For Canada, the chart shows Alberta's Heritage Fund and Alberta's GDP.

Table 4
Predictive regressions, one month ahead

	Panel A				With cont. S&P returns			Panel B				With cont. S&P returns		
	Country	oil^a_{t-1}	oil^b_{t-1}	oil^c_{t-1}	oil^a_{t-1}	oil^b_{t-1}	oil^c_{t-1}	Country	oil^a_{t-1}	oil^b_{t-1}	oil^c_{t-1}	oil^a_{t-1}	oil^b_{-1}	oil^c_{-1}
slope	Canada	0.0809	0.0763	0.0809	0.0610	0.0503	0.0438	Nigeria	0.1441	0.1213	0.1144	0.1372	0.1123	0.1009
t-stat		2.01	2.13	2.49	2.46	2.10	1.95		1.88	1.68	1.72	1.96	1.67	1.62
R^2_{adj}		0.015	0.013	0.017					0.023	0.016	0.015			
slope	Colombia	0.0639	0.0702	0.0917	0.0511	0.0536	0.0682	Norway	0.0647	0.0643	0.0819	0.0424	0.0352	0.0403
t-stat		1.05	1.21	1.70	0.90	0.96	1.28		1.22	1.32	1.79	1.07	0.95	1.19
R^2_{adj}		0.001	0.002	0.007					0.003	0.003	0.008			
slope	Ecuador	0.1074	0.0948	0.1232	0.1092	0.0973	0.1278	Oman	0.1374	0.1149	0.1099	0.1301	0.1066	0.0970
t-stat		1.18	1.14	1.41	1.20	1.18	1.46		3.04	2.96	3.28	3.28	2.98	2.93
R^2_{adj}		0.009	0.007	0.015					0.061	0.043	0.044			
slope	Egypt	0.0465	0.0157	0.0028	0.0328	−0.0023	−0.0231	Russia	0.1793	0.2076	0.1821	0.1476	0.1663	0.1231
t-stat		0.71	0.25	0.04	0.60	−0.04	−0.42		1.91	2.39	2.22	1.88	2.21	1.70
R^2_{adj}		−0.001	−0.003	−0.004					0.011	0.016	0.013			
slope	Indonesia	0.1699	0.1625	0.1450	0.1479	0.134	0.1042	Saudi Arabia	0.1251	0.1204	0.1374	0.1165	0.1091	0.1227
t-stat		2.02	2.03	1.98	2.03	1.89	1.61		2.27	2.13	2.47	2.33	2.13	2.44
R^2_{adj}		0.015	0.014	0.011					0.027	0.025	0.037			
slope	Kuwait	0.1987	0.1655	0.1447	0.1478	0.1165	0.0866	Tunisia	0.0553	0.0577	0.0527	0.0547	0.0569	0.0516
t-stat		3.18	2.95	2.91	2.69	2.28	1.83		1.95	2.02	1.94	1.96	2.02	1.89
R^2_{adj}		0.092	0.067	0.056					0.010	0.011	0.009			
slope	Malaysia	0.0494	0.057	0.0662	0.0358	0.0393	0.0409							
t-stat		0.91	1.07	1.34	0.72	0.79	0.88							
R^2_{adj}		0.000	0.002	0.004										
slope	Mexico	0.0361	0.0471	0.0428	0.0129	0.0168	−0.0008							
t-stat		0.69	1.01	1.00	0.39	0.54	−0.03							
R^2_{adj}		−0.002	0.000	−0.001										

The table shows coefficients, *t*-stats, and adjusted R^2s from regressions of monthly excess stock returns on lagged changes in the price of WTI oil. The three rightmost columns of each panel show results from regressions of monthly excess stock returns on lagged changes in the price of WTI oil and contemporaneous excess returns on the S&P 500 index. The superscripts *a*, *b*, and *c* indicate that the daily oil series is lagged by one, two, and three days, respectively, before calculating monthly oil price changes. The regressions include one lag of the monthly oil series at the time.

quickly than they incorporate sectoral news about oil prices (see Driesprong et al., 2008, Hong et al., 2007, and Fan and Jahan-Parvar (2012) for a thorough discussion). As a result, the predictive relation reflects sectoral economic news (including oil price changes) rather than changes in risk premia.

The bulk of our results are based on the following predictive regression:

$$r_t^i = \alpha + \beta_0 r_t^{S\&P} + \beta_1 oil_{t-1} + \epsilon_t^i, \tag{1}$$

where r_t^i is the log return on country *i*'s stock market index in month *t* and in excess of the riskless rate, $r_t^{S\&P}$ is the excess log return on the S&P 500 index in month *t*, and oil_{t-1} is the log change in oil prices in month $t-1$. Eq. (1) can be interpreted as a variant of international CAPM (ICAPM), with the S&P 500 as the global factor. We use the S&P 500 since it is highly correlated with and a substantial constituent of all global equity indexes.[6] Unless speci-

[6] Bear in mind that: 1) As noted by Kilian and Park (2009), real crude oil prices are largely subject to the same shocks that impact the U.S. economy, and, by extension, the U.S. stock market. 2) We use price-dividend ratios as a robustness measure. They are readily available for S&P 500 index, but not for global measures such as MSCI ACWI. 3) The United States equities account for a substantial portion of global equity index construction. For example, large cap U.S. stocks constitute 53% of MSCI ACWI index. As a result, broad U.S. equity indexes such as S&P 500 and global indexes such

Table 5
Predictive regressions, including three lags of oil changes

	Country	oil^a_{t-1}	oil^a_{t-2}	oil^a_{t-3}	With cont. S&P returns oil^a_{t-1}	With cont. S&P returns oil^a_{t-2}	With cont. S&P returns oil^a_{t-3}
slope	Canada	0.0751	0.0519	0.0268	0.0610	0.0064	−0.0143
t-stat		1.89	1.25	0.76	2.50	0.28	−0.64
slope	Indonesia	0.1601	0.0747	0.0822	0.1447	0.0252	0.0375
t-stat		1.91	0.89	1.09	1.97	0.34	0.59
slope	Kuwait	0.1451	0.1080	0.0581	0.1020	0.0825	0.0429
t-stat		2.63	1.86	1.07	2.11	1.46	0.99
slope	Nigeria	0.1286	0.1295	0.092	0.1246	0.1143	0.0775
t-stat		1.94	1.99	1.16	1.99	1.76	1.01
slope	Oman	0.1297	0.0649	0.0606	0.1269	0.0431	0.0460
t-stat		2.99	1.92	1.75	3.16	1.35	1.35
slope	Russia	0.1698	0.1133	0.0543	0.1473	0.0406	−0.0114
t-stat		1.82	1.38	0.63	1.86	0.56	−0.15
slope	Saudi Arabia	0.1175	0.0604	0.0340	0.1106	0.0426	0.0091
t-stat		2.18	1.32	0.80	2.23	1.00	0.23
slope	Tunisia	0.0549	0.0164	−0.0118	0.0545	0.0149	−0.0131
t-stat		1.89	0.61	−0.38	1.90	0.56	−0.44

The table shows coefficients and *t*-stats from regressions of monthly excess stock returns on up to three lags of changes in the price of WTI oil. The three rightmost columns show results from regressions of monthly excess stock returns on lagged changes in the price of WTI oil and contemporaneous excess returns on the S&P 500 index. The superscript *a* indicates that the daily oil series is lagged by one day before calculating monthly oil price changes.

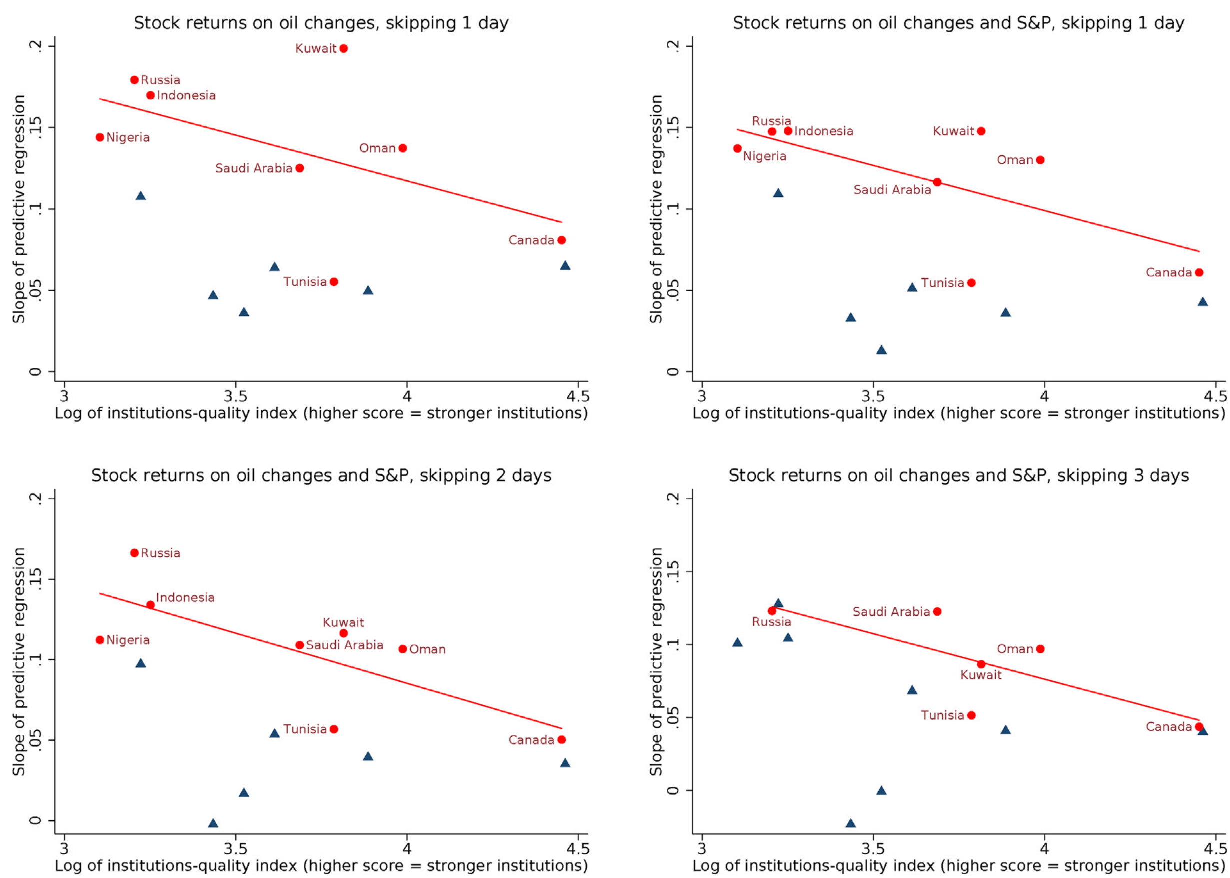

Fig. 2. Corruption and stock return predictability. For each country, the charts plot the slopes of predictive regressions against the corresponding log-IQIs. Red markers show countries for which the slopes are statistically significant. Each chart indicates whether contemporaneous S&P 500 excess returns are included, and the number of days by which the oil time series is lagged.

fied otherwise, local stock returns include dividend reinvestment, and oil price changes are calculated using WTI prices. The riskless rate is the appropriately compounded 3-month Treasury Bill rate (series DTB3 from the Federal Reserve Economic Date website). The coefficients are estimated with ordinary least squares, and statistical significance is evaluated using heteroskedasticity-consistent standard errors.

The regressions include the contemporaneous return on the S&P 500 index to account for broad market movements, so to minimize the risk that oil price changes reflect economic news about the U.S., which is the largest oil consumer during our sampling period. In some of the robustness checks, we include the S&P 500

as MSCI ACWI are highly correlated. 4) Our results are generally robust to inclusion of a measure such as MSCI ACWI returns.

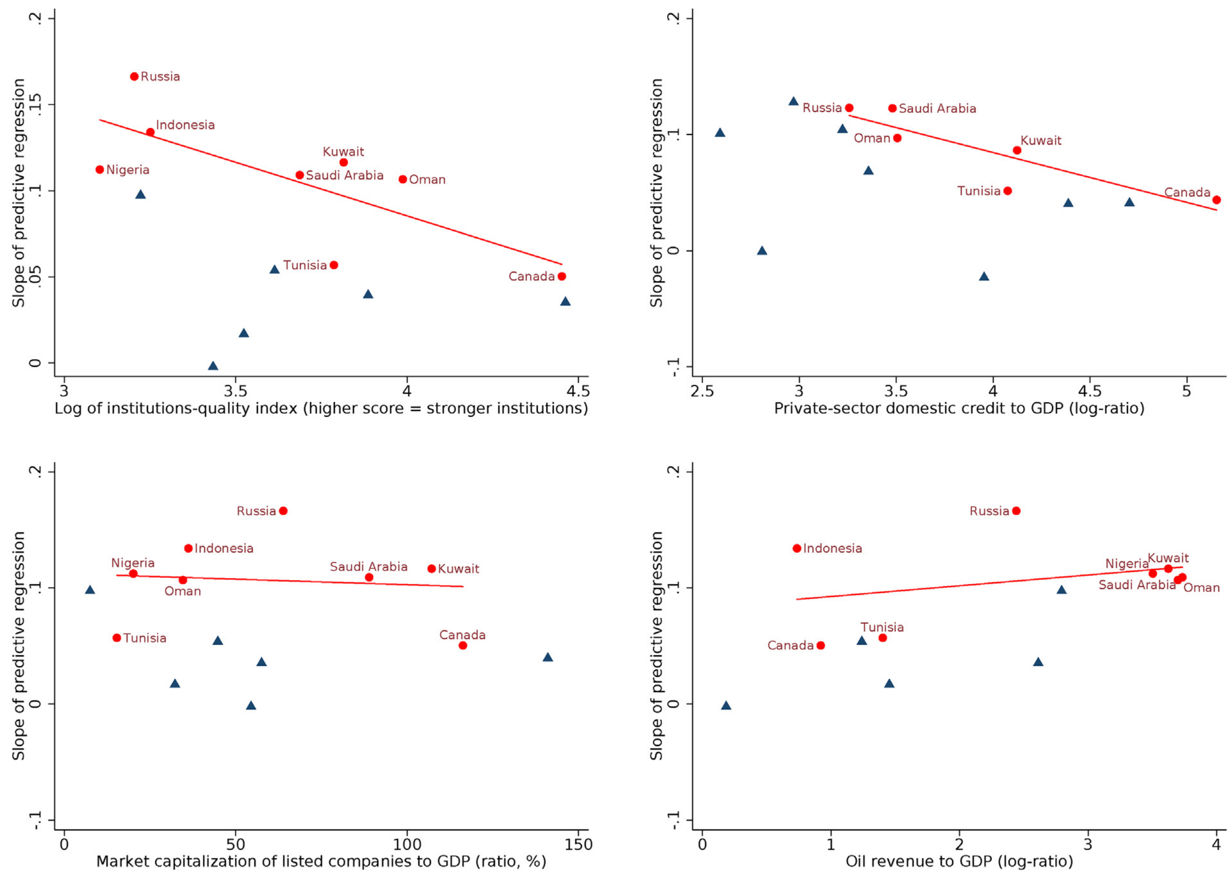

Fig. 3. Corruption index and stock return predictability: Possible alternative explanations. For each country, the charts plot the slopes of predictive regressions against the corresponding log-IQIs, private-sector domestic credit to GDP, capitalization of listed companies to GDP, and oil revenue to GDP. Red markers show countries for which the slopes are statistically significant. Regressions include contemporaneous S&P 500 excess returns and daily oil prices are lagged by two days.

log dividend-price ratio as a predictor to understand whether oil price changes are proxies for valuation ratios (both S&P 500 returns and the dividend-price ratio are from the dataset used by Goyal and Welch, 2008, and available on Amit Goyal's website).[7] The log dividend-price ratio is highly autocorrelated, which can bias predictive regression coefficients (Stambaugh, 1999). As a consequence, we use the procedure in Amihud, Hurvich, and Wang (2009) to obtain both the coefficients and the standard errors.

Our analysis is predominantly based on a set of country-specific regressions. Our main results are not based on panel regressions for the following reasons. First, our primary interest is to evaluate the response of equity market returns to oil price changes separately for each country. Second, we let the slope coefficients of contemporaneous S&P 500 returns vary across countries to account for varying degrees of correlation between the local and U.S. stock markets. Using a plain vanilla panel estimator would average both the oil and S&P 500 coefficients across countries, and introducing dummy variables to differentiate country effects on both the oil and S&P 500 slope parameters is tantamount to running a set of country-specific regressions. However, as a robustness measure, we fit the data with the random-coefficient panel regression model of Swamy (1970). This method allows us to account for parameter heterogeneity across individual countries. Effectively, the estimated model is characterized as:

$$y_{i,t} = X_{i,t}\beta_i + \epsilon_{i,t}, \tag{2}$$

where i indexes the various countries and $X_{i,t}$ is a set of appropriately lagged regressors including oil price changes, S&P 500 returns, the S&P 500 price-dividend ratio, and squared oil price changes. In this specification, the parameter vector β_i is the coefficient for the i^{th} country, such that:

$$\beta_i = \beta + \nu_i, \qquad E(\nu_i) = 0, \qquad E(\nu_i \nu_i') = \Sigma_\nu. \tag{3}$$

As a result, the estimated model is:

$$y_{i,t} = X_{i,t}(\beta + \nu_i) + \epsilon_{i,t} = X_{i,t}\beta + (X_{i,t}\nu_i + \epsilon_{i,t}). \tag{4}$$

Additional details can be found in Greene (2012).

In our first set of regressions, we lag the oil price series by one day before calculating monthly oil price changes. As a result, the change in month t is computed from the second-to-last day in month $t-1$ to the second-to-last day in month t. The reason is that we compare predictability in markets across a wide range of time zones, from Malaysia to Canada. Without lagging the oil time series, oil price news on a given day could be reflected in Canadian stock prices on the same day, and in Malaysian stock prices the following day. As a consequence, cross-country differences in predictability could be due to asynchronous trading.

In some results we also lag the oil time series by two or three days before calculating oil price changes; in these cases we are using lags to evaluate how long it takes for information in oil prices to be reflected in stock prices, in line with the analysis in Driesprong et al. (2008). We use a superscript to indicate by how many days the oil series is lagged: oil^a (one day), oil^b (two days), and oil^c (three days).

[7] See, among many other studies, Campbell and Shiller (1988), Fama and French (1988), Campbell and Thompson (2008), and Goyal and Welch (2008).

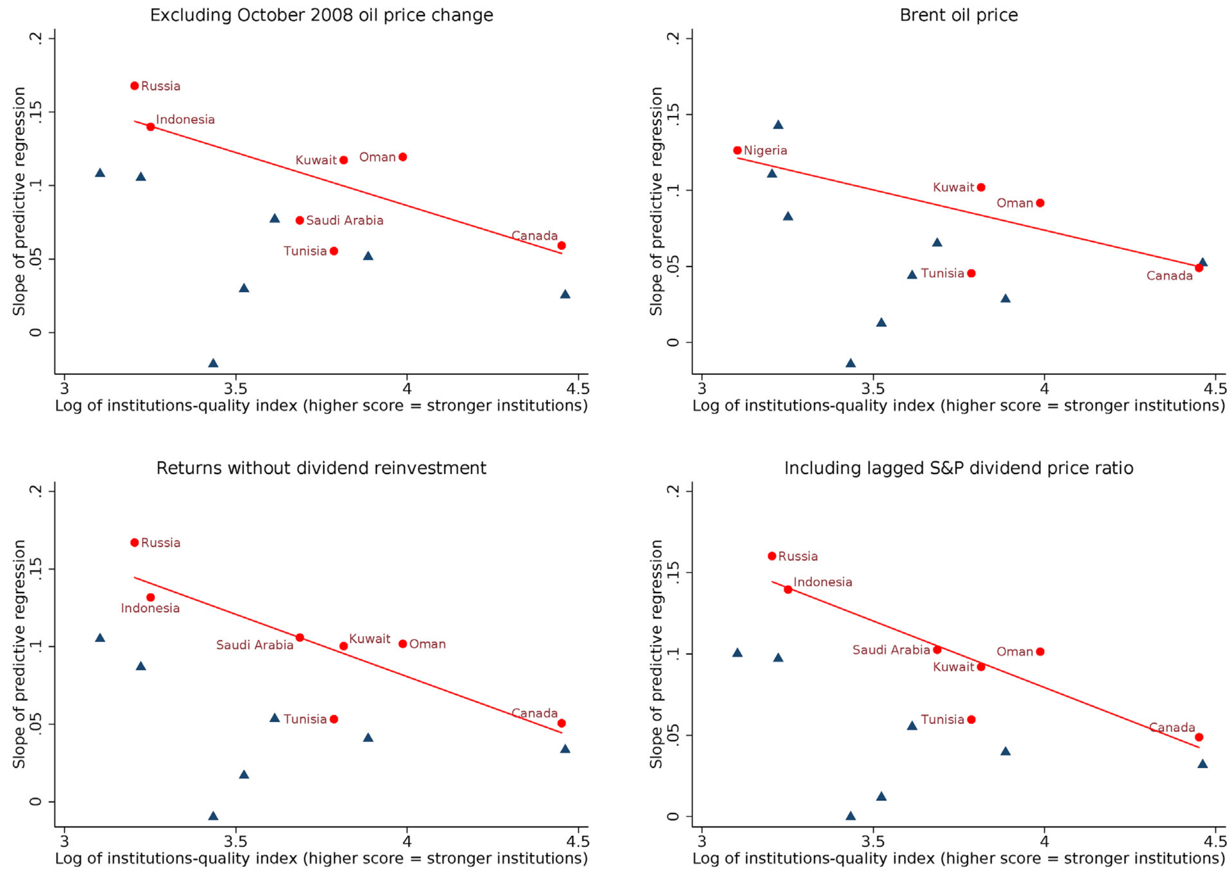

Fig. 4. Corruption index and stock return predictability: Robustness checks. For each country, the charts plot the slopes of predictive regressions against the corresponding log-IQIs when (1) excluding the October 2008 oil price change; (2) using Brent oil prices instead of WTI oil prices; (3) returns exclude dividends; (4) replacing contemporaneous S&P 500 returns with the S&P 500 lagged price-dividend yield as a regressor. Red markers show countries for which the slopes are statistically significant. Unless specified otherwise, regressions include contemporaneous S&P 500 excess returns and daily oil prices are lagged by two days.

3. Results

As shown in Table 4, we find statistically significant predictability for eight of the 14 countries we study. The first three columns of each panel show slopes and t-stats for regressions in which only lagged oil price changes are included as predictors. The slow diffusion of information theory states that new information from a market (in our case, crude spot market) takes a while to be impounded into prices in another market (in our case, the equity market). How long this process of slow diffusion lasts is an empirical question. We address it by following Driesprong et al. (2008) and Fan and Jahan-Parvar (2012) and experiment with skipping a finite number of days in the previous month ($t-1$) while building the lagged oil price change measure. The predictive strength, statistical significance, and fit generally decline as the number of days by which the oil series are lagged increases, although for all the eight countries the predictive slopes are still statistically significant at 10% when lagging the oil series by three days (note that in all cases we regress one-month ahead equity returns on monthly oil price changes; the daily lags indicate how many days we skip between the oil price change in month $t-1$ and the equity return in month t). With the exception of Kuwait and Oman, the adjusted R^2s are fairly small, which is typical for predictive regressions and is due to the volatility of returns. The patterns are largely unchanged when including contemporaneous S&P 500 excess returns in the regressions, with the oil slopes (β_1 in Eq. (1)) slightly smaller, and the statistical significance higher. These results support the hypothesis that oil prices reflect, in part, information about the state of the U.S. economy.

Our findings are noticeably different from those of Driesprong et al. (2008, Table 5), who provide evidence of no predictability for the eight countries that, in our analysis, have predictable returns. The discrepancy is driven by samples that overlap only partially. We find no predictability when we run the predictability regressions with samples ending in April 2003, which is when the sample of Driesprong et al. (2008) ends.

Driesprong et al. (2008) conclude that the predictive power of oil is not driven by risk premia, but by delays in information processing. Our empirical findings confirm Driesprong et al.'s conclusions. We also observe that predictability is driven by delays in information diffusion. In Table 5 we report predictive slopes and t-stats from regressions of returns on three lags of oil price changes for the eight countries that have significant slopes in Table 4. We expect to find no longer-term predictive power if predictability is driven by slow information diffusion rather than changes in risk premia associated with oil price fluctuations. For five countries, only the first lag of oil price changes has predictive power. In the case of Kuwait, the second lag has a statistically significant slope, while for Oman it is the third lag that is weakly statistically significant. Overall, these results suggest that predictability is due to frictions in the diffusion of oil price information rather than changes in risk premia, as documented by Driesprong et al. (2008).

We now turn to how the quality of a country's institutions affects the predictive regression, and we should point out that our analysis going forward is largely based on a visual representation of the results. The reason is that our sample selection criteria – discussed in Section 2.1 – limit the sample to 14 countries, which constrains

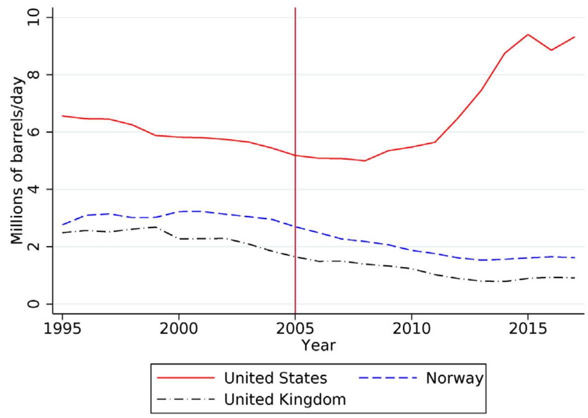

Fig. 5. Crude oil production in the U.S.A., U.K., and Norway. The chart shows the time series of oil production in the USA, UK, and Norway, from 1980 to 2017. Production is measured as the average daily number of barrels, in millions. The vertical line highlights the year 2005.

our ability to carry out a formal statistical analysis in the cross-section.

In Fig. 2 we show scatter plots of the slopes of predictive regressions (vertical axis) against the log IQI, with red markers indicating countries for which the predictive relation is statistically significant (the absolute value of the *t*-stat being greater than 1.65).[8] We use the log of the index because Canada and Norway have a much higher score than the other countries, and they would exert a strong leverage effect if the index were not transformed. In the top left chart, the slopes are computed from regressions that lag the oil price series by one day and that do not include the contemporaneous return on the S&P 500. In the remaining regressions, the contemporaneous return on the S&P 500 is included, and the oil price series is lagged by one, two, and three days, as indicated.

All charts imply a negative relationship between institutional quality and the size of the predictive slope coefficient, even without differentiating on the basis of statistical significance. When considering statistically significant slopes only, lower institutional quality is clearly associated with stronger predictability. Moreover, the lack of statistical significance is not related to the level of institutional quality.

We argue that the absence of predictability in some of the countries shown in Table 4 and also in Fig. 2 is likely to be driven by the low liquidity of the respective equity markets. Pastor and Stambaugh (2003) study the asset-pricing implications of liquidity, which they define as "the ability to trade large quantities [of assets] quickly, at low cost, and without moving the price". In countries with less developed financial markets, one may observe relatively large price movements that reflect a lack of liquidity rather than changes in fundamentals. As a consequence, it may be more difficult to establish a relation between stock returns and lagged oil prices.

The rightmost columns of Table 1 show the ratios of market capitalization of listed companies to GDP (which we take as a measure of financial development) in 1995, 2003, and 2010. Using 2003 data to proxy for the level throughout the sample period, the eight countries for which we find statistically significant predictability have an average ratio of 53%, compared to 44% for countries with statistically insignificant predictive slopes. The median ratios, which reduce the impact of Malaysia's outsized value, are 39% and 25%, respectively, indicating that countries for which we find no predictability have substantially less developed financial markets. We should emphasize that, as discussed below, the pattern of predictive power across financial development does not match the pattern of predictive power across institutional quality very well. Overall, these results suggest a threshold effect of financial development: identifying predictability is difficult for low levels of financial development, but higher levels are not associated with the strength of the predictive relation.

In Fig. 3 we evaluate whether the relation between predictability and institutional quality can be explained by leverage, financial development, or the importance of oil exports relative to the economy.[9] To save space, we only show results when the oil price series is lagged by two days. The top left chart, which is taken from Fig. 2 to facilitate the comparison with the remaining charts, shows IQI values against predictive slopes calculated when including contemporaneous S&P 500 returns.[10] The first alternative explanation we consider is leverage: if companies in countries with low institutional quality have higher leverage, their stock prices would mechanically be more reactive to changes in oil prices, and the predictive slopes would be larger. In the top right chart of Fig. 3, we plot the slopes against the log-ratio of private-sector domestic credit to GDP. We can see that higher predictive power is actually associated with lower leverage, the opposite of what we would expect if the predictability/institutional quality pattern was due to leverage. The fact that the two top charts look very similar is consistent with the observation that legal investor protection is positively related with the development of financial markets, including debt markets (La Porta et al., 1997).

The bottom left chart of Fig. 3 shows the predictive slopes against the capitalization of listed companies relative to GDP, which we consider a measure of financial development. Russia, which is the country with the largest predictive slope, has a fairly developed stock market. Tunisia, on the other hand, ranks lowest in financial development even though predictability is about as strong as for Canada, which has the largest stock market relative to GDP. The predictive slope for Oman is about the same as in the case of Kuwait and Saudi Arabia, even though Oman's stock market is relatively small.

The last chart in Fig. 3 evaluates the effect of oil-export dependence on the predictive relation. It could be the case, for instance, that a larger predictive slope simply reflects that a country derives a larger amount of revenue from exporting oil. The bottom right chart shows that the Middle-Eastern countries export significantly more oil than the rest in relation to GDP, but predictability is not commensurately stronger for these countries relative to Russia or Indonesia. Indonesia derives the smallest amount of revenue from oil exports, yet it has the second strongest predictability below Russia and just above Kuwait.

We provide an additional set of robustness checks in Fig. 4. First, we exclude the October 2008 oil price change from the sample (top left). Second, we use Brent oil prices instead of WTI prices (top right). Third, we consider returns without dividend reinvestment (bottom left). Finally, we include the lagged S&P 500 price dividend ratio as a regressor in place of the contemporaneous S&P 500 return. Oil prices experienced their largest negative return in October 2008, and we exclude that month's observation to evaluate whether it exerts an undue leverage on the results. The predictive relation

[8] We follow the asset pricing literature in focusing on slope parameters of predictive regressions instead of, for example, R^2 measures. See Fama and French (1989) and Campbell and Yogo (2006).

[9] Our measure for leverage is private-sector domestic credit to GDP (series FS.AST.PRVT.GD.ZS from the World Bank). The values shown are averages between 2003 and 2013. Financial development and oil exports are defined in Section 2.1.

[10] In the Figures, we show average values for the various country-specific variables. The IQI data is only available from 2003 for all the countries we study, and, for comparability reasons, averages for the other variables are also calculated starting from 2003 to 2012, the year through which when most variables are available. In the case of private-sector domestic credit to GDP, data is only available through 2006.

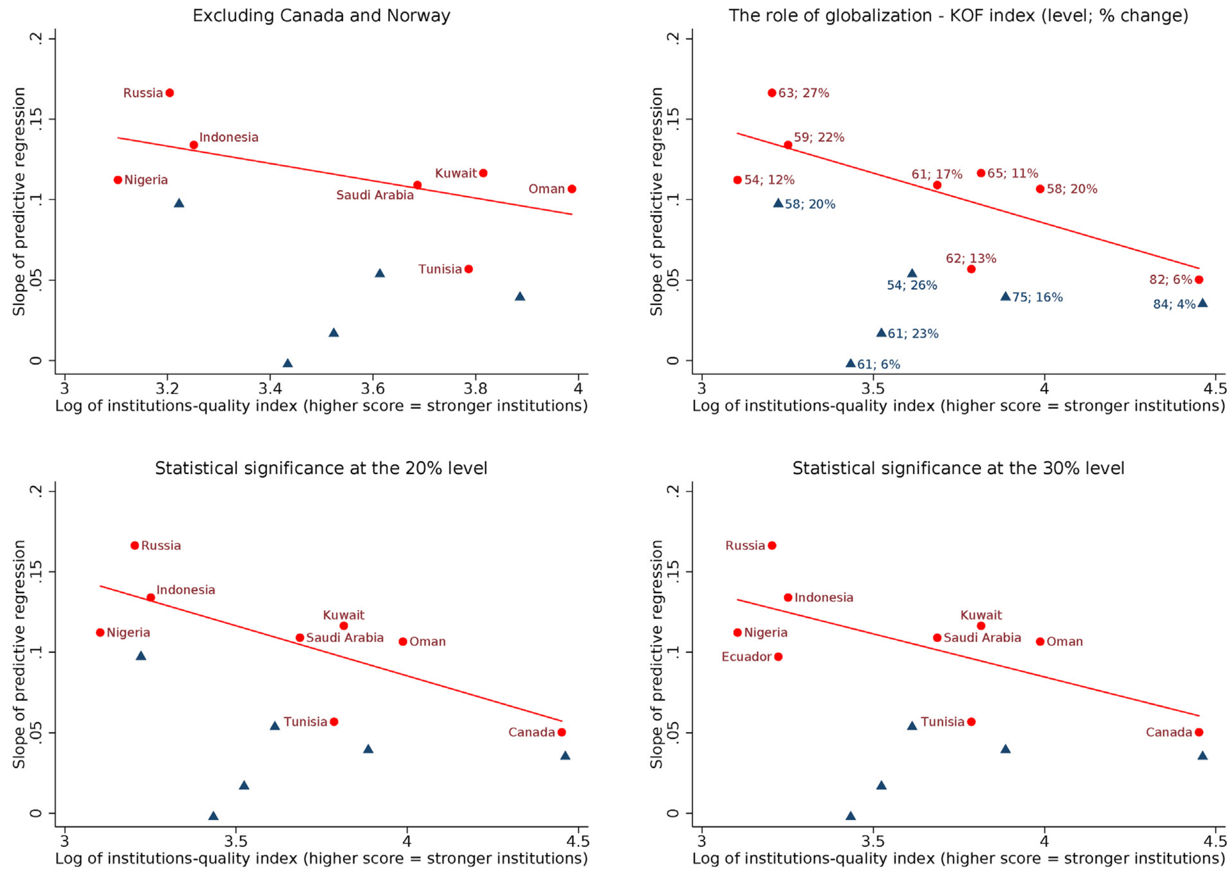

Fig. 6. Corruption index and stock return predictability: Additional robustness checks. The top left chart shows the slopes of predictive regressions against the corresponding log-IQIs when excluding Canada and Norway, whose relatively high IQI values might exert a leverage effect. In the top right chart, we consider the effect of globalization by plotting the slope of predictive regression against the corresponding log-IQI. We report the average value and the change in the ETH Zurich KOF Globalization Index (See Dreher (2006)). Gygli, Haelg, and Sturm (2018) over the sample for each country. In the bottom charts, we consider a slope coefficient to be statistically significant if the absolute value of the *t*-statistic is above 1.28 and 1.04, respectively. These values correspond to statistical significance at the 20% and 30% levels. Red markers show countries for which the slopes are statistically significant.

becomes statistically insignificant for Saudi Arabia, but it remains broadly unchanged for the other countries. We should highlight that, while we need to be aware of the potential effect of large oil price movements on statistical inference, these movements generate substantial variation in the countries' available revenue and thus provide identification. As shown in the bottom panels of Fig. 4, excluding dividends from the calculation of returns and including the lagged log dividend-price ratio does not alter our predictability results.

Predictability is weaker when using Brent prices (from the Federal Reserve Economic Data) instead of WTI prices, with the slopes becoming statistically insignificant for Indonesia and Saudi Arabia, and the relation between IQI and the slopes being less steep overall. While price changes for Brent and WTI oil are highly correlated, the correlation varies over time. It is equal to 89% between 1995 and 2013, but it is higher early in the sample and it peaks at 96% between 2002 and 2004. It is lower in the second half of the sample, with a minimum of 70% between 2011 and 2013. This change in price correlation coincides with a drastic change in production levels in the United States (WTI is a US benchmark price) relative to Norway and the United Kingdom (Brent is a North-Sea benchmark). Fig. 5 shows the time series of the crude oil production between 1995 and 2014, with 2005 highlighted with a vertical red line. While production in Norway and the United Kingdom continued the decline that started around 2000 and dropped by 42% and 52%, respectively, between 2005 and 2014, production in the United States increased by 69% over the same years. The divergence in volume likely makes the price of WTI crude more informative than Brent crude in the latter part of the sample, which is an important period for identifying the relation between oil price and stock returns.

Continuing with robustness checks, we report the estimation results after the exclusion of Canada and Norway in the top left chart of Fig. 6. The IQI scores for these two countries are noticeably higher than for the rest of our sample, and we want to ascertain whether they exert an unduly large leverage effect on the trend line. We observe that such is not the case. The trend line remains practically the same with and without the two countries. We then turn on the effect of globalization on our findings. We use the KOF Globalization index, maintained by the KOF Swiss Economic Institute at ETH Zurich. We calculate the average index value, and the change in the index over our sample, for each of the 15 countries we study. We then show the average and change next to each country's observation in the baseline plot of predictive coefficients and quality of institutions on the top right panel of Fig. 6. We do not observe a systematic pattern in how the globalization index relates to the predictive coefficients. In general, the countries we study tend to be relatively open, with all but two of them having an index average above the world average in the sample period (56). In the bottom panels of Table 6, we evaluate the effect of using alternative thresholds for statistical significance. They report the estimation results when using 20% and 30%, respectively, as the significance levels. The results carry through virtually unchanged.

In Fig. 7, we evaluate whether the predictive relation between stock returns and oil price changes may include a non-linear term.

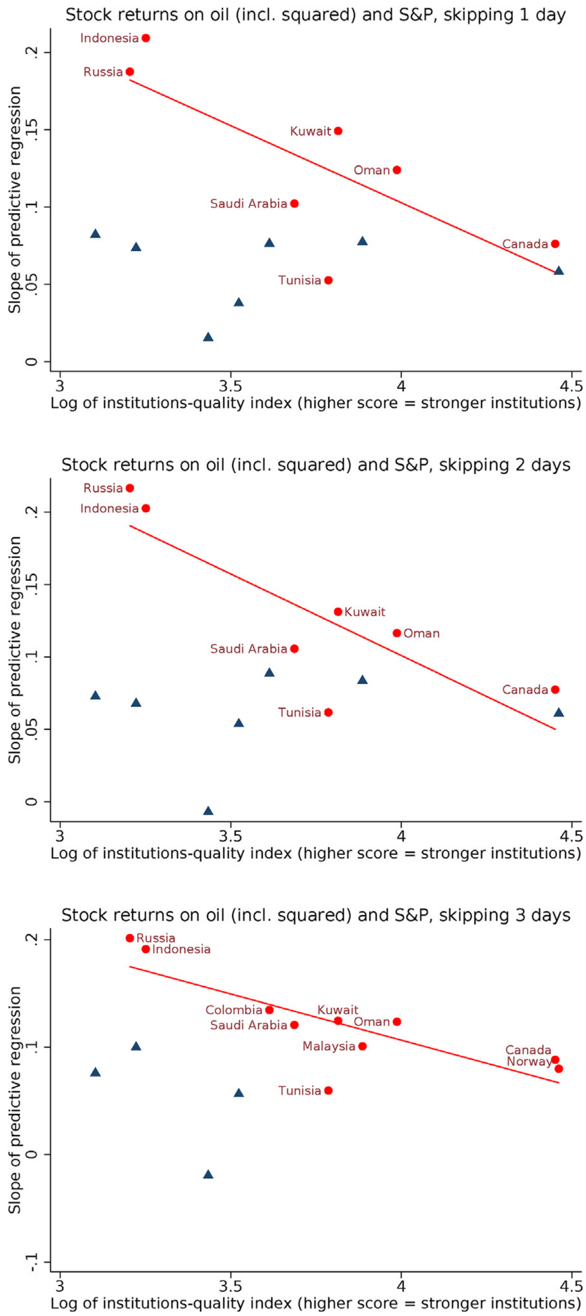

Fig. 7. Corruption index and stock return predictability: Including squared lags of oil price changes. For each country, the charts plot the slopes of predictive regressions against the corresponding log-IQIs when lagged squared oil price changes are included along with oil price changes. The three charts show results when the daily oil price series is lagged by (top to bottom) one, two, or three days before the monthly changes are calculated. Red markers show countries for which the slopes are statistically significant.

We then include the lagged squared value of oil price changes together with the linear component, whereas the daily oil price series is lagged by one, two, or three days before monthly oil changes are computed. A comparison of the three panels in Fig. 7 with the top left panel of Fig. 2 reveals that the results are largely unchanged.

The effect of oil price changes could be different depending on whether the price of oil increases or decreases. Investors tend to be more sensitive to downside risk, hence in Fig. 8 we evaluate the predictive power of negative net oil price changes. These changes are defined following Hamilton (1996), as the log-difference of spot

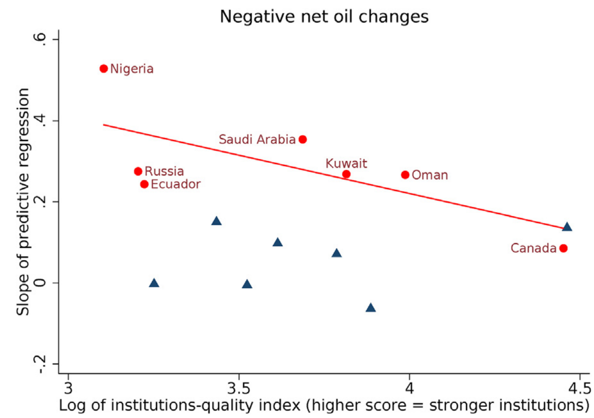

Fig. 8. Corruption index and stock return predictability: Negative oil price changes. For each country, the chart plots the slopes of predictive regressions against the corresponding log-IQIs when using negative net oil price changes (Hamilton, 1996). Red markers show countries for which the slopes are statistically significant.

oil price at time *t* and the minimum of the spot oil price over the previous three months, if the current spot oil price is below this minimum, and zero otherwise. The predictive coefficients of negative net oil price changes clearly line up against the quality of institutions as expected, and the magnitude of the coefficients is about three times as large as in our baseline results.

As mentioned in Section 2.2, we estimate Swamy (1970)'s random coefficient panel regression model to account for cross-sectional parameter heterogeneity. The estimated slope parameters from this exercise, plotted against the log-IQI, are shown in Fig. 9. In the left-hand side panels, we include lagged oil changes as the only regressor, while in the right-hand side panels we include both lagged oil changes and squared lagged oil changes as regressors. In the top panels, the daily oil price series is lagged by one day before the monthly changes are computed, while in the bottom panels the series is lagged by two days. As shown in the figure, the results are remarkably similar to those presented so far. In addition, Colombia, Ecuador, and Malaysia have statistically significant coefficients in at least some of the charts. Note that, to account for potential cross-sectional correlation in the response of stock returns to oil price changes, we consider a coefficient to be statistically significant if the absolute value of its *t*-statistic is greater than 1.96, rather than 1.65.

3.1. Institutions, oil prices, and equity values

The results discussed above can be used to approximate the magnitude of the impact of institutional quality on how oil revenue is spent. We do so by taking the difference between predictive slopes for countries with high and low institutional quality. The top left chart of Fig. 3, which we take as representative of the results, shows that the difference between the interpolated statistically significant slopes at the highest and lowest IQI levels is slightly greater than 0.1. If we round this difference down to 0.1, a 1% increase in oil prices translates into 0.1% higher stock market returns for countries with low institutional quality relative to those with high institutional quality.

Note that the slope coefficients are informative about the impact of oil price changes on equity prices, not on asset prices. The reason is that equity-holders have a claim on the firm's assets which is subordinated to that of debt-holders. In order to calculate the effect of oil price changes on the value of the assets held by listed companies we need information on the liabilities issued by the firms included in the stock market indexes, which is unavailable for most countries. We can, however, use a formula that describes the relation

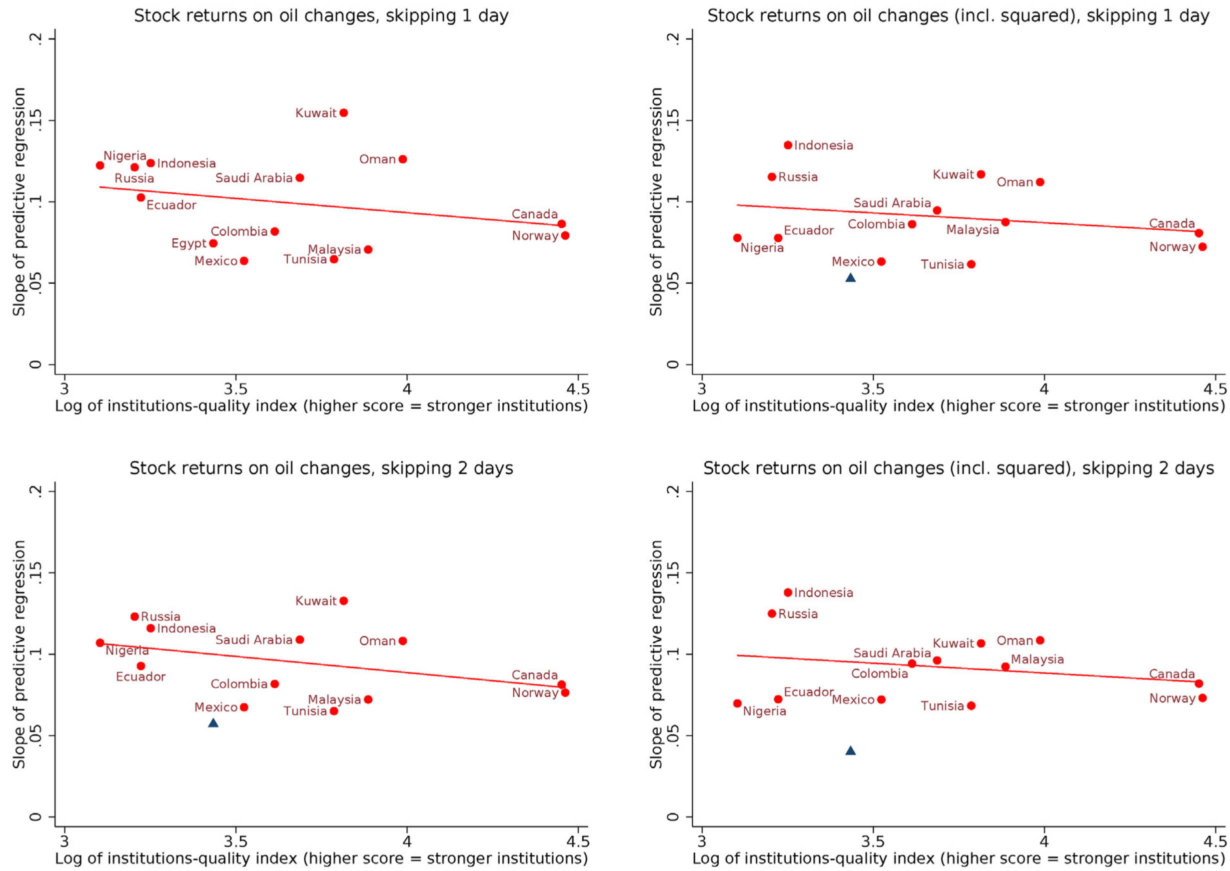

Fig. 9. Corruption index and stock return predictability: Results based on random-coefficient panel regressions. For each country, the charts plot the slopes of predictive regressions against the corresponding log-IQIs when the daily oil price series is lagged by one day (top panels) or two days (bottom panels) before the monthly changes are calculated, and the regressions include only oil price changes (left panels) or oil price changes and squared oil price changes (right panels). Red markers show countries for which the slopes are statistically significant.

between the slope coefficients of equity and asset returns to estimate the effect of oil price changes on asset values, for different leverage levels.[11] For a ratio of debt to equity equal to 1 (or 3), a 1% increase in oil prices would result in asset values 0.050% (or 0.025%) higher in countries with low institutional quality relative to those with high institutional quality.

The pattern of access to debt markets across countries, as shown in Fig. 3, makes our results conservative. The reason is that based on this figure, countries with weaker institutions have lower leverage and – as discussed above – the sensitivity of stock returns to oil price changes depends on the leverage of the companies in the stock index: the higher the leverage, the higher the sensitivity. If their leverage were increased to match that of countries with stronger institutions, their slope coefficients would also increase, and the spread in the sensitivity of returns to oil prices between low- and high-institutional quality countries would be larger.

We should emphasize that our intention is not to accurately measure how institutional quality affects fiscal policy, but to gauge the approximate magnitude of its implications on the expenditure of oil revenue using a simple methodology that is applicable to panels of commodity-exporting countries. Below we discuss the main limitations of our analysis, which originate from data constraints.

In some countries, the market capitalization of listed companies is relatively low compared to GDP. The Canadian stock market, for instance, is about three times as large as its Indonesian counterpart (see Table 1). Such difference raises the question of whether the effect of oil price on equity returns, as measured for listed companies, can be generalized to non-listed companies, and whether the ability to generalize varies significantly across countries. Cross-country variation is especially important to us because our conclusions are based on comparing countries across the spectrum of institutional quality.

The first source of cross-country variation is that listed companies in a country may not be representative of the financing decisions of all firms – listed and non-listed – in that economy. The difference lies in leverage, since small companies in countries with more developed financial systems may have better access to debt than their counterparts in countries with less developed financial markets. A study of the differences in financing decisions between listed and non-listed firms across countries is an interesting question. Unfortunately, due to unavailability of data, it cannot be addressed in our study.

Second, there could be a compositional effect, because listed companies could be disproportionately more active in oil-related sectors in countries for which oil exports are more important. However, predictability does not appear to be related to the size of oil exports to GDP (Fig. 3). For instance, Indonesia (which has the second-strongest predictability) has an oil export/GDP ratio similar to Canada (whose returns are the least predictable).

[11] Let D and E be the amount of debt and equity issued by a firm. Also, let β_1^E be the predictive slope calculated using stock returns. The predictive slope that is applicable to asset values can be calculated from leverage (D/E) and β_1^E as follows: $\beta_1^A = \frac{\beta_1^E}{1+D/E}$.

4. Conclusions

We study whether stock market returns in oil-exporting countries can be predicted by oil price changes, and we investigate the link between predictability and the quality of each country's institutions, which we measure with perceived corruption levels. Returns are predictable for about half the countries we consider, and predictability is stronger when institutional quality is lower.

We argue that the relation between predictability and institutional quality reflects the preference of countries with weaker institutions to consume oil windfalls locally through pro-cyclical fiscal policies, rather than smooth out the impact of windfalls by, for instance, investing the proceeds abroad. Within this framework, our results can be used to gauge the approximate magnitude of institutional quality's impact on how oil revenue is spent.

References

Alesina, A., Tabellini, G., & Campante, F. (2008). Why is fiscal policy often procyclical? *Journal of the European Economic Association*, *6*(5), 1006–1036.

Amihud, Y., Hurvich, C., & Wang, Y. (2009). Multiple-predictor regressions: Hypothesis testing. *Review of Financial Studies*, *22*(1), 413–434.

Barsky, R., & Kilian, L. (2004). Oil and the macroeconomy since the 1970. *Journal of Economic Perspectives*, *18*(4), 115–134.

Campbell, J., & Shiller, R. (1988). The dividend-price ratio and expectations of future dividends and discount factors. *Review of Financial Studies*, *1*(3), 195–228.

Campbell, J., & Thompson, S. (2008). Predicting excess stock returns our of sample: Can anything beat the historical average? *Review of Financial Studies*, *21*(4), 1509–1531.

Campbell, J. Y., & Yogo, M. (2006). Efficient tests of stock return predictability. *Journal of Financial Economics*, *81*(1), 27–60.

Dreher, A. (2006). Does globalization affect growth? Evidence from a new index of globalization. *Applied Economics*, *38*(10), 1091–1110.

Driesprong, G., Jacobsen, B., & Maat, B. (2008). Striking oil: Another puzzle? *Journal of Financial Economics*, *89*(2), 307–327.

Esfahani, H., Mohaddes, K., & Pesaran, M. (2014). An empirical growth model for major oil exporters. *Journal of Applied Econometrics*, *29*(1), 1–21.

Fama, E., & French, K. (1988). Dividend yields and expected stock returns. *Journal of Financial Economics*, *22*(1), 3–25.

Fama, E., & French, K. (1989). Business conditions and expected returns on stocks and bonds. *Journal of Financial Economics*, *25*, 23–49.

Fan, Q., & Jahan-Parvar, M. R. (2012). U.S. industry-level returns and oil prices. *International Review of Economics & Finance*, *22*(1), 112–128.

Fatás, A., & Mihov, I. (2003). The case for restricting fiscal policy discretion. *Quarterly Journal of Economics*, *118*(4), 1419–1447.

Fisman, R., & Miguel, E. (2007). Corruption, norms, and legal enforcement: Evidence from diplomatic parking tickets. *Journal of Political Economy*, *115*(6), 1020–1048.

Gogineni, S. (2010). Oil and the stock market: An industry level analysis. *Financial Review*, *45*(4), 995–1010.

Goyal, A., & Welch, I. (2008). A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies*, *21*(4), 1455–1508.

Greene, W. (2012). *Econometric Analysis* (7th edn). New Jersey: Prentice Hall.

Gygli, S., Haelg, F., & Sturm, J. (2018). *The KOF globalisation index - Revisited*. Working paper.

Hamilton, J. (1983). Oil and the macroeconomy since World War II. *Journal of Political Economy*, *91*(2), 228–248.

Hamilton, J. (1996). This is what happened to the oil price-macroeconomy relationship. *Journal of Monetary Economics*, *38*(2), 215–220.

Hamilton, J. (2011). Nonlinearities and the macroeconomic effects of oil prices. *Macroeconomic Dynamics*, *15*(S3), 364–378.

Hong, H., & Stein, J. C. (1999). A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets. *Journal of Finance*, *54*(6), 2143–2184.

Hong, H., Torous, W., & Valkanov, R. (2007). Do industries lead stock markets? *Journal of Financial Economics*, *83*(2), 367–396.

Huang, R., Masulis, R., & Stoll, H. (1996). Energy shocks and financial markets. *Journal of Futures Markets*, *16*(1), 1–27.

James, A. (2015). The resource curse: A statistical mirage? *Journal of Development Economics*, *114*, 55–63.

Jones, C., & Kaul, G. (1996). Oil and the stock markets. *Journal of Finance*, *51*(2), 463–491.

Kaufmann, D., Kraay, A., & Mastruzzi, M. (2005). Governance matters IV: Governance indicators for 1996–2004. In *Policy Research Working Paper no. 3630*.

Kaufmann, D., Kraay, A., & Mastruzzi, M. (2006). Governance matters V: Governance indicators for 1996–2005. In *Policy Research Working Paper no. 4012*.

Kilian, L. (2008). The economic effects of energy price shocks. *Journal of Economic Literature*, *46*(4), 871–909.

Kilian, L., & Park, C. (2009). The impact of oil price shocks on the U.S. stock market. *International Economic Review*, *50*(4), 1267–1287.

Kotter, J., & Lel, U. (2011). Friends or foes? Target selection decisions of sovereign wealth funds and their consequences. *Journal of Financial Economics*, *101*(2), 360–381.

La Porta, R., Lopez-De-Silanes, F., Shleifer, A., & Vishny, R. (1997). Legal determinants of external finance. *Journal of Finance*, *52*(3), 1131–1150.

Leite, C., & Weidmann, J. (1999). Does Moher Nature corrupt? Natural resources, corruption, and economic growth. *IMF Working Paper*.

Mauro, P. (1995). Corruption and growth. *Quarterly Journal of Economics*, *110*(3), 681–712.

Mehlum, H., Moene, K., & Torvik, R. (2006). Institutions and the resource curse. *Economic Journal*, *116*(508), 1–20.

Pastor, L., & Stambaugh, R. (2003). Liquidity risk and expected stock returns. *Journal of Political Economy*, *111*(3), 642–685.

Persson, T. (2002). Do political institutions shape economic policy? *Econometrica*, *70*(3), 883–905.

Pieschacon, A. (2012). The value of fiscal discipline for oil-exporting countries. *Journal of Monetary Economics*, *59*(3), 250–268.

Reboredo, J., & Rivera-Castro, M. (2014). Wavelet-based evidence on the impact of oil prices on stock returns. *International Review of Economics & Finance*, *29*, 145–176.

Sachs, J., & Werner, A. (1995). Natural resource abundance and economic growth. *NBER Working Paper 5398*.

Stambaugh, R. (1999). Predictive regressions. *Journal of Financial Economics*, *54*(3), 375–421.

Swamy, P. (1970). Efficient inference in a random coefficient regression model. *Econometrica*, *38*, 311–323.

van der Ploeg, F. (2011). Natural resources: Curse or blessing? *Journal of Economic Literature*, *49*(2), 366–420.

Vicente, P. (2010). Does oil corrupt? Evidence from a natural experiment in West Africa. *Journal of Development Economics*, *92*(1), 28–38.

Wilhelm, P. G. (2002). International Validation of the Corruption Perceptions Index: Implications for Business Ethics and Entrepreneurship Education. *Journal of Business Ethics*, *35*(3), 177–189.