Copenhagen Business School

Master Thesis

MSc Economics and Business Administration

Finance and Investments

The Time-Varying Relationship between Oil Prices and Stock Returns: Evidence from Oil-Importing and Oil-Exporting Countries

Authors:

Francesca Giraudo (S133344)

Ilse Aldegonda Wilhelmina Hutten (S133073)

Supervisor:

Kasper Lund-Jensen

Standard Pages: 76

Characters: 157,145

Submission Date: May 17, 2021

Abstract

This paper studies the relationship between oil prices and stock markets of four oil-importing and six oil-exporting countries. Monthly data is used for the period March 1983- December 2020. First, a rolling-window regression model is used to establish the time-varying contemporaneous relationship between WTI futures prices and the stock indices. The industrial production index is included as a variable to control for global growth. Second, we test whether investors react with a delay to oil price changes as hypothesised by Hong and Stein (1999). Our findings suggest that there is a significant relationship for all countries at least at some point in time. The oil-exporting countries generally show a significantly positive relationship, consistent with the theoretical expectations. The oil-importing countries show less homogeneous results. The results of our second test provide evidence for underreaction behaviour of investors, as the oil effect strengthens once we include lags of several trading days between monthly oil prices and stock returns. Moreover, we demonstrate that trading strategies based on the oil effect generate superior gains in comparison with buy-and-hold strategy.

Table of Contents

1. Introduction

On April 20th, 2020, the price of crude oil became negative for the first time in history. Oil prices plummeted below zero and kept decreasing, with the West Texas Intermediate (WTI) crude oil price tumbling to a negative $37 per barrel (Tobben, 2020). The reason behind this drastic decrease in oil prices is essentially the mismatch between oil supply and demand. The Covid-19 pandemic has caused an unprecedented drop in demand for petroleum as governments across the world put into force lockdowns, forcing citizens at home and restricting international travel. Meanwhile, the major oil-producers continued to produce increasing amounts of oil, causing the oil prices to decrease dramatically. This incidence of negative oil prices had a strong impact on stock markets of both oil-importing and exporting countries, highlighting the importance of understanding the returns caused by such event shocks for investors to reduce risks when investing.

Even though oil prices had never been negative before, oil shocks have continuously affected the real economy and stock markets throughout time. Oil prices affect the stock markets of both oil-importing and exporting countries through their influences on monetary policy instruments, inflation, corporate income and other economic activities (Gourène & Mendy, 2018). Understanding this effect on international stock markets is fundamental for market participants to make informed decisions. Therefore, a bulk of existing literature studied this relationship and most of the studies agree that there exists one. There is consensus in the literature that a rise in oil prices increases stock prices in oil-exporting countries, mainly due to higher revenues from oil exports (Kilian & Park, 2009). The findings regarding oil-importing countries are mixed (i.e., Sadorsky, 1999; Papapetrou, 2001; Narayan & Narayan, 2010). A limited number of studies finds no statistically significant relationship (i.e., Apergis & Miller, 2009; Miller & Ratti, 2009).

With these premises, we believe it is of interest to further examine the relationship between oil prices and stock markets. First, we analyse whether there exists a contemporaneous relationship between oil price changes and stock market returns. Second, we analyse whether this relationship strengthens after several trading days, as investors often react with a delay.

For the first part of the analysis, we aim to investigate whether oil prices have an effect on the stock markets of various countries and how this effect evolves over time. We consider six primarily oil-exporting countries (Canada, Russia, Norway, Brazil, Saudi Arabia and

Mexico) and four primarily oil-importing countries (United States, China, Jordan and Hong Kong) during the period March 1983 - December 2020. To model the time-varying relationship between oil prices and stock markets, a rolling-window regression is adopted with the addition of the industrial production factor. The industrial production index serves as a control variable that attempts to control for global growth. Our findings confirm the existence of a contemporaneous relationship between oil prices and stock returns. We find heterogeneous results depending on the specific period in time and the nature of the country (oil-importing versus oil-exporting). For all oil-exporting countries, besides Saudi Arabia, we find a positive and almost entirely significant relationship. For the oil-importing countries, the relationship is observed to be less homogeneous and less significant. We attribute these heterogeneous findings to the notion that the oil market is subject to exogenous forces and to the asymmetric relationship between oil prices and the stock market.

For the second part of the analysis, we are interested in analysing whether the oil effect is stronger when investors' underreaction is taken into consideration. Investors' underreaction is often observed, as it appears to take time before information is fully reflected in the stock market (Hong & Stein, 1999). Investors do not react in a timely manner to new information as they are subject to several behavioural biases or because they find it difficult to correctly assess the impact of such information on stock prices. In order to test whether this delayed reaction exists in the oil-stock market relationship, a static regression model is used where lags of several trading days are added between the monthly oil prices and stock returns. We compare the strength (i.e., explanatory power) of the oil effect for the different daily lags and find evidence that confirms a delayed reaction of investors. More precisely, for the majority of the countries, the explanatory power spikes at day eight, implying that it takes around eight trading days before information about oil prices is fully incorporated in the market. Exceptions are Jordan and Russia, where the oil effect is the strongest at lag zero, thereby signalling little underreaction. We believe it is of interest to examine if investors react with a delay to oil prices as such underreaction challenges the notion of market efficiency and can create short-term predictability. Therefore, investors can possibly exploit this market inefficiency to predict stock prices and obtain abnormal returns.

To demonstrate the applicability of the established relationship in a real-world setting, we develop two trading strategies based on the oil effect, while taking investor underreaction into consideration. Both strategies are compared to a simple buy-and-hold strategy. By evaluating the risk and return features of the three strategies, we find that the two oil strategies overall outperform the buy-and-hold one.

Existing research on the relationship between oil prices and stock returns provides mixed results, where no full consensus on the directionality of this relationship has been reached. Therefore, we believe it is of interest to further investigate the dynamics of this relationship. This thesis documents several contributions to this field of research. First, unlike other studies, this paper takes the time-varying relationship into consideration. Second, a diverse set of countries is chosen to examine the effect of oil prices on both oil-exporting and importing, as well as developed and emerging countries. Third, this paper assesses the relationship over an extended period, that includes the most recent recession caused by the Covid-19 pandemic, therefore providing insights on current events. Fourth, this paper investigates not only the purely theoretical relationship between oil prices and stock markets, but also takes investors' behaviour into account by testing for underreaction. Combining behavioural finance with traditional finance allows us to holistically analyse the relationship.

The remainder of this paper is structured as follows. Section 2 provides a brief overview of the existing empirical literature on the topic. Section 3 discusses the theoretical literature that lays the foundation for our analysis. Section 4 presents the two hypotheses that are tested in the paper. Section 5 describes the series of data employed in the analysis. Section 6 outlines the econometric model used in the empirical analysis to test both hypotheses. Section 7 presents the results of the analysis. Section 8 discusses our findings, in light of our expectations and previous research. Section 9 outlines the limitations of our analysis and provides suggestions for future research. Section 10 presents the trading strategies developed based on oil price changes and discusses the main implications. Finally, Section 11 concludes by highlighting the main findings of our analysis.

2. Literature Review

Despite oil being one of the world's most influential commodities, there appears to be limited research conducted on the possible linkages between oil and equity markets. The initial lack of research in the field can be explained by the fact that oil prices have been relatively stable up until 1973. The oil crisis in 1973, however, caused the oil prices to fluctuate and triggered the emergence of a new body of literature interested in analysing oil prices and their impact on the economy. The primary focus of the early empirical analysis was the possible linkage between oil and macro-economic factors such as Gross Domestic Product (GDP), inflation and interest rates. Later, in the 1990s, commodity-based financial assets were developed, leading to an increase in trading in the commodity (futures) markets. This process, referred to as the financialization of commodity markets, led to a rise in empirical research examining the direct relationship between oil prices and equity markets (Cheng & Xiong, 2013).

To better explain market anomalies, a new branch of literature, known as behavioural finance, arose in the 1980s. Behavioural finance studies the behaviour of investors, by researching how emotional and psychological aspects can influence market outcomes. Several studies within our topic take a behavioural approach, as they argue that behavioural biases influence the dynamics of the oil-stock market relationship.

In the following section, the literature relevant for our paper is reviewed. Section 2.1 presents the main studies that investigate the relationship between oil prices and stock markets. The empirical research on both aggregate and disaggregate level is discussed. Section 2.2 presents oil-stock market research that is more consistent with behavioural approaches, as well as research that introduces models to explain this behaviour.

2.1 The relationship between oil and stock markets

The first influential paper is the work of Hamilton (1983), who provides evidence of a link between oil and the macro-economy. He argues that the majority of the recessions in the United States from World War II until the early 1980's has been anticipated by sharp increases in the prices of crude petroleum, thus suggesting a significant negative relationship with the economic activity. However, the study does not touch upon the possible connections between stock market and oil prices. His work is followed by a bulk of research that studies the impact of oil prices on various macroeconomic variables (i.e., Burbidge & Harrison, 1984; Loungani, 1986, Gisser & Goodwin, 1986).

Prior to the late 1990s, the number of studies investigating whether oil price changes directly impact stock market returns is rather limited. Some exceptions are the studies conducted by Chen et al. (1986) and Jones and Kaul (1996). Chen et al. (1986) research whether several macroeconomic variables and oil prices systematically affect stock returns. They find that industrial production and interest rates have a statistically significant impact on stock returns. However, as far as crude oil is concerned, they find that the oil coefficients are insignificant over their sample period of 1968-1977 and that oil price risk is not priced in stock markets. Jones and Kaul (1996) test whether the reaction of four stock markets, Canada, the United Kingdom, Japan and the United States, to oil shocks can be justified by current and future changes in real cash flows and/or changes in expected returns, using the standard cash flow dividend valuation model. The results show that for the United States and Canada this reaction can be accounted for entirely by the impact of the oil shocks on cash flows, showing rationality in the equity markets. However, the model used does not yield strong evidence for Japan and the United Kingdom, whose markets show excess volatility. The authors suggest that oil price changes might forecast stock market returns but leave the issue for further research.

After the late 1990s, an increasing number of studies that specifically examines the impact of oil prices on stock prices was conducted, on both aggregate and disaggregate level. The results of the empirical research on aggregate level are mixed, as papers report both negative and positive as well as inconclusive results. A negative relationship is found by, among others, Sadorsky (1999) who researches the relationship between oil prices and stock return for the U.S. equity market with a vector autoregression (VAR) model during the period 1947-1996. The findings show that oil prices and their volatility have a strong impact on stocks returns. In particular, the analysis reveals that the equity returns decrease in the short-term in response to a rise in oil prices. These results are in line with Papapetrou (2001), who studies the relationship among oil prices, stock returns, interest rates, economic activity and employment in Greece during the period 1989-1999, using a vector error model. The evidence shows that oil price changes affect economic activity and employment, and that they have a negative effect on stock returns during the first four months. Further, Driesprong et al. (2008) investigate whether changes in oil prices can be used to predict stock returns of 48 countries, both developed and emerging, during the period 1973-2003. By applying a linear regression model, they find that higher oil prices predict lower stock returns. They argue that this negative relationship cannot be attributed to time-varying risk premia. However, they suggest that the evidence is more consistent with the idea that investors react at different points in time to changes in oil prices. Another negative relationship is found by Asteriou and Bashmakova (2013), who assess the impact of oil prices on emerging stock markets, namely Central and Eastern European Countries. A panel data approach is used for the period 1999-2007. The findings indicate a negative reaction of stock returns to upward and downward movements of the oil market, with the results being highly significant in case of low oil prices. These studies are in line with the majority of the research, which argues that an oil price increase leads to higher costs and thus lower stock returns.

On the contrary, a positive relationship is found by Narayan and Narayan (2010), who examine the relation between stock prices and oil prices for Vietnam over the period 2000-2008. They include the exchange rate as an additional determinant of stock prices and find that equity prices, oil prices and nominal exchange rates are cointegrated. Both oil prices and exchange rates have a statistically significant positive effect on stock prices in the long run. Reasons behind this positive link could be that the Vietnamese stock market has grown impressively in the last years, triggered by internal and domestic factors. Therefore, it is likely that the impact of these factors on the Vietnamese equity market was more dominant than the oil price increase. This is in line with Kilian (2008a) who states that positive deviations in oil prices caused by unexpected global economic growth positively affect stock returns within the first year of the expansionary shock. The essence of this paper is that the effect of a rise in oil prices is dependent on its underlying cause.

Alternatively, some studies find insignificant results. An example is the empirical analysis of Huang et al. (1996), who investigate the relationship between daily oil futures returns and daily U.S. stock returns using a VAR approach. They find that oil futures returns do impact the stock returns of some individual oil companies, but no significant relation is detected between oil futures returns and general market indices, such as the S&P 500. Apergis and Miller (2009) examine whether structural oil-market shocks affect stock returns in eight developed countries during the period 1981–2007. They document no significant responses of international stock markets to oil price shocks. Similarly, Maghyereh (2004) examines the link between oil price shocks and equity returns for 22 emerging countries by applying a VAR model over the period 1998-2004. The findings show that there is no significant evidence that crude oil prices have an impact on stock index returns in these countries. He concludes that, unlike previous empirical studies in developed economies, stock markets in emerging economies are inefficient in the transmission of new information regarding the oil market, and stock market returns in these countries do not rationally signal changes in crude oil price. Further, Cong et al. (2008) investigate the effects of oil prices changes on Chinese stock market returns with a multivariate vector autoregression. They find no statistically significant impact on the stock returns of most Chinese stock market indices, except for the manufacturing index and several oil companies. Last, Miller and Ratti (2009) analyse the long-run relationship between oil prices and international stock markets over the period 1971-2008. The oil price effect measured over the entire period does not yield any statistically significant results. They do find short-term evidence of a negative relationship between the oil price and six OECD [1] countries for the periods 01/1971 to 05/1980, and 02/1988 to 09/1999, respectively. After 09/1999, the long-run relationships are again mostly insignificant, except for Canada, for which they find a positive link. Their findings support a conjecture of change in the relationship between oil prices and stock prices in the last decade, which may suggest the presence of stock market or oil price bubbles.

The negative relationship established by the majority of the literature on aggregate level does not necessarily hold for the stock markets in oil-exporting countries, according to among others, Mohanty and Nandha (2011). They assess the relation between changes in oil prices and equity returns in Gulf Cooperation Council (GCC) countries using country-level data. Except for Kuwait, stock markets have significant positive exposures to oil price changes, showing that the aforementioned effects on stock markets do not apply to all countries. This is in line with Wang et al. (2013), who investigate the oil-price relationship in the major oil-importing and exporting countries with a VAR analysis. They find that positive aggregate demand shocks show a higher degree of co-movement among the stock markets in oil-exporting countries compared to those in oil-importing countries. Therefore, they conclude that the response of a country's stock market to oil price changes is highly dependent on whether the country is a net oil-importer or oil-exporter in the aggregate demand. This conclusion also holds for Park and Ratti (2008), whose analysis focuses on the United States and 13 European countries. They find that the impacts of oil price changes on the Norwegian stock market are positive, while the stock markets of oil-importing countries are negatively affected. In addition, Antonakakis et al.

[1] OECD stands for "Organisation for Economic Co-operation and Development".

(2013) examines the influence of oil prices on stock markets over five stock market indices from both oil-importing (United States, United Kingdom and Germany) and oil-exporting countries (Canada and Norway). They again find that the oil effect depends on whether the country is an importer or exporter, but also that the correlation is not constant over time (i.e., time-varying).

The aforementioned studies research the oil-stock relationship on aggregate level. However, researching on aggregate level may mask heterogeneous responses from different industrial sectors due to their different characteristics and responses to variations in oil prices. These characteristics are related to whether the sectors are classified as oil-users, oil-substitutes or non-oil-related. Most sector-level research is country specific. For example, Sadorsky (2001) uses a multi-factor market model to estimate the expected returns to the Canadian oil & gas industry stock prices. An increase in the market factor or oil price factor increases the return to Canadian oil and gas companies. This is in line with El-Sharif et al. (2005) who analyse oil & gas returns in the United Kingdom. In addition, they note that non-oil & gas sectors are weakly linked to oil price changes. Studies with a multi-country setting on industry level are rather limited. Two noteworthy studies have been conducted by Nandha and Faff (2008) and Arouri and Nguyen (2010). Nandha and Faff (2008) examine the linkages between oil price shocks and equity pricing using the standard market model with an oil factor included in the regression. They apply the model to 35 global industry indexes for the period 1983-2005. They find that oil price rises have a negative impact on equity returns for all sectors, except for the mining and oil & gas industries. Arouri and Nguyen (2010) conduct a sector-level analysis in Europe using the multi-factor asset pricing model and the Granger causality test. They obtain a strong and significant correlation between oil and stock prices in most European countries and conclude that the reactions of stock returns to oil price changes differ greatly depending on the activity sector. More precisely, the oil & gas industry shows a high positive sensitivity to oil price changes, whereas the food & beverage sector presents a negative sensitivity to oil price changes. Overall, the results on sector-level are more consistent across studies. That is, the studies provide strong evidence that oil-related and oil-substitute sectors, such as the oil & gas and mining sector, are positively affected by changes in oil prices, whereas oil-user and non-oil-related sectors are negatively affected.

2.2 Underreaction to oil price changes in stock markets

Several studies regarding the oil-stock market relationship take a behavioural approach, as they argue that behavioural biases influence the dynamics of this relationship. For instance, Driesprong et al. (2008), Fan and Jahan-Parvar (2011) and Pollet (2011) find that oil prices can predict stock market returns. However, they argue that this predictability is not consistent with the notion of time-varying risk premium and offer behavioural finance explanations that can possibly cause predictability.

First, Driesprong et al. (2008) state that the predictability they find based on oil prices cannot be attributed to the time-varying risk premium. Their results show that higher oil prices (also) predict lower stock returns, whereas assuming a time-varying risk premium, high oil prices would lead to higher economic risk and thus higher returns, and vice versa. Thus, they suggest that this predictability can be explained by the underreaction behaviour of investors, as it takes time before information about oil price changes is fully reflected in the stock market. They research whether the predictability of the stock market is strengthened when lags of several trading days are added between monthly stock index returns and monthly lagged oil price changes. For the majority of the countries in their sample, they find that the explanatory power of the regressions increases up to a lag of six trading days and then starts decreasing. Similarly, Fan and Jahan-Parvar (2011) investigate whether investors react with a delay to information regarding oil price changes. They conduct an industry-level research on the United States and find a delayed response of up to two trading weeks to news about oil prices. The results are significant for all sectors and the underreaction is more pronounced in non-oil related sectors, such as business services and construction. They argue that investors active in such sectors are more likely to underestimate the economic effect of oil price changes and fail to react in a timely manner. Pollet (2011) also researches underreaction to oil price changes in several industries in the United States and finds significant underreaction in non-oil dependent sectors, such as construction, household products, medical equipment, real estate, and financial services.

The delayed reaction of investors found in the papers of Driesprong et al. (2008), Fan and Jahan-Parvar (2011) and Pollet (2011) can be explained by the gradual diffusion hypothesis. This hypothesis has been presented by Hong and Stein (1999), who develop a dynamic model of a single asset in which private information diffuses gradually across the investing public. The investors are unable to extract information from prices. Thus, underreaction in stock prices to news occur, resulting in short-term stock return stability and stock price momentum. Hong et al. (2007) extend the framework of the gradual information diffusion model to include publicly available information. They argue that stock returns might still observe underreaction even though the information itself is publicly available (i.e., oil prices), as opposed to privately available. Their investigation is based on the idea that investors are only boundedly rational, that is, they are only able to process small subsets of the available information in an unbiased way. They find that gradual information diffusion can lead to underreaction (and thus cross-asset return predictability) when investors do not pay attention to news from other markets, react to information at different points in time, or have difficulties in assessing the impact of public information on the value of stocks. The empirical findings discussed in this section confirm that the gradual diffusion hypothesis holds not only for private information, but also for public information as suggested by Hong et al. (2007).

2.3 Relevance

Our study adds to the existing literature in several ways. First, the existing studies on the effect of oil price changes on stock market returns yield mixed results, namely positive, negative and insignificant relationships are found, depending on the scope of the countries and the period covered. Thus, the lack of consensus asks for further research in this area. Second, the majority of previous research does not account for the time-varying relationship between oil price changes and stock market returns. This paper uses a methodology that takes this time-varying aspect into consideration. That is, a regression with a rolling-window is employed to capture how the relationship changes over time. Third, most studies are focused on the United States, which is primarily an oil-importer. As the impact of oil price changes on the equity markets is likely to differ depending on whether a country is an oil-importer or oil-exporter, it is essential to consider both oil-importing and oil-exporting countries. Fourth, studies who focus on countries other than the United States, are mainly considering developed countries. The effect of oil price changes is possibly different on the financial markets of developing countries versus emerging countries. Consequently, this research includes a diverse set of countries, both oil-importing and oil-exporting, as well as developed and emerging, to conduct an extensive analysis of oil price effects. Fifth, this paper assesses the relationship over a relatively long period compared to other papers, that spans from March 1983 to December 2020. This time span is characterized by periods of high volatility and instability of the financial markets, especially during the Global Financial Crisis in 2008/2009 and the most recent Covid-19 recession. To conclude, understanding the relationship between oil prices and stock markets, while taking underreaction behaviour into account, is particularly relevant for investors as this knowledge can possibly be exploited in a trading strategy with the purpose of earning excess returns.

3. Theoretical Review

This section discusses the theory that is relevant for understanding how oil prices can impact stock returns and how oil prices can potentially create predictability. Section 3.1 covers the theoretical transmission mechanisms by which oil price changes can affect stock markets. In Section 3.2 and 3.3, several views on the predictability of stock prices are discussed. First, the Efficient Market Hypothesis (EMH) and the Random Walk Hypothesis (RWH) are explained. Thereafter, behavioural finance theories and their implications are introduced, with an emphasis on the theory of underreaction, as this theory will be tested later in the empirical analysis.

3.1. Transmission Channels between Oil Prices and Stock Prices

In this section, the theoretical transmission mechanisms by which oil price changes can affect stock markets are explained. The main identified transmission channels affecting stock markets are the stock valuation, monetary, output, fiscal and uncertainty channel.

3.1.1 Stock valuation channel

The stock valuation channel is the direct transmission mechanism by which oil prices influence a company's cash flows and thus the stock prices. Economic theory suggests that current stock prices reflect the discounted future cash flows of a particular stock (Williams, 1938). This can be shown as:

$$P_{i,t} = \sum_{i=1}^{n} E\left(\frac{CF_n}{(1+r)^n}\right) \tag{1}$$

where CF_n is the cash flow at time n, E is the expectation operator and r is the discount rate. Therefore, any factor affecting the discounted future cash flow has a significant impact on the stock price and consequently on the stock return. Oil prices can affect a firm's future cash flows either negatively or positively, depending on whether the firm is an oil-consumer or oil-producer. For an oil-consuming firm, oil is a production factor, implying that an increase in oil prices results in an increase in production costs. An increase in costs reduces the profit and thus future cash flows, negatively affecting stock prices and returns. The opposite holds true for oil-producers that benefit from an oil price increase. This event leads to a growth in profits and improves future cash flows, thus positively influencing stock prices.

3.1.2 Monetary channel

Another transmission mechanism by which oil price changes affect stock returns is the monetary channel, that is, through inflation and interest rates. Namely, oil prices affect the discount rate, also referred to as the required rate of return, which intuitively reflects the time value of money and the risk premium demanded by investors. The discount rate is at least partly composed of the expected inflation and expected real interest rates (Mohanty & Nandha, 2011). Thus, by affecting inflation and interest rates, oil price changes consequently influence stock returns.

Theoretically, on a firm level, a rising oil price results in increased production costs and thus higher prices for consumers. At a national level, these higher oil prices can lead to inflation as the general price level rises (Hamilton, 1996). As the goal of monetary policy makers is to maintain a low and stable inflation rate, they typically respond to an increase in inflation with an increase in short-term interest rates (Cologni & Manera, 2008). Higher short-term interest rates drive to higher commercial borrowing rates for future investments. Due to these increased borrowing costs, firms have fewer positive net present value projects, resulting in a reduction in future cash flows. Hence, the response from central banks to a positive deviation in oil price is likely to increase the short-term interest rates, thus resulting in a decrease in cash flows and an overall reduction in stock price value.

3.1.3 Output channel

The output channel represents the mechanism through which oil price fluctuations affect aggregate output. An increase in oil price impacts aggregate output through both a production cost effect and an income effect. As the production cost effect is already presented in Section 3.1.1, the focus of this section is the income effect. The income effect entails that, for an oil-importing country, an increase in oil prices causes a reduction in the discretionary income of households, due to a rise of retail prices, as well as gasoline and heating oil prices (Bernanke, 1983). Lower income leads to lower consumption and thus aggregate output, which further diminishes labour demand. Stock markets have the tendency to respond negatively to such developments. In particular, lower aggregate demand leads to lower expected cash flows for firms, which further causes a reduction in stock prices, as seen in Equation 1. Conversely, oil-exporting countries will experience a positive income shock when the oil prices increase. These countries will benefit from increased oil revenues as the value of export demand for oil rises, leading to higher aggregate demand and thus higher output. The positive change in aggregate demand will, however, only occur if the positive income effect is stronger than the negative production cost effect.

3.1.4 Fiscal channel

The fiscal channel is mainly applicable to oil-exporting countries that are financing their physical and social infrastructure through their oil revenues (Ayadi, 2005). A surge in oil prices increases the wealth of an oil-exporting country, hence allowing for increased government spending. In case private consumption and government spending are complements of each other, then government spending will lead to higher household consumption and thus higher expected cash flows for firms. The opposite occurs due to the crowding out effect in case private consumption and government spending are substitutes. The crowding out effect refers to the paradox that an expansionary fiscal policy can drive down private sector spending (Cologni & Manera, 2008). Stock markets will respond negatively to such developments. To sum up, stock returns in an oil-exporting country can either be positively or negatively impacted via the fiscal transmission channel, depending on whether private consumption and government spending are complements or substitutes.

3.1.5 Uncertainty channel

The last identified transmission channel is the uncertainty channel. Oil price fluctuations induce higher uncertainty in the real economy, resulting among others from shifts in inflation, aggregate demand and output (Brown & Yucel, 2002). An increase in uncertainty can reduce firms' demand for investments as they cannot foresee whether the increase in energy prices will be transitory or long-lasting. A delay in investments dampens the prospects of economic growth and stock returns. Thereby, high uncertainty can also affect the risk premium demanded by investors as a compensation for the risk they take. Last, in times of uncertainty, households tend to reduce their consumption of durable goods. There is an increased incentive to save instead of to consume, which also has a negative impact on the economy and stock market returns (Edelstein & Kilian, 2009).

3.1.6 Summary of the channels

The five transmission channels explain the mixed results of the literature overview. The combination of effects within a specific channel or across different channels can generate ambiguous responses from the stock markets. For instance, oil prices can affect a firm's future cash flows either positively or negatively, depending on whether the firm is an oil-consumer or oil-producer. Thereby, at a national level, higher oil prices can lead to inflation as the general price level rises. Central banks typically respond to an increase in inflation with an increase in interest rates. Hence, a rise in oil price is likely to increase the discount rates and consequently decrease the cash flows, resulting in an overall reduction of stock price value. In addition, a rise in oil prices generates a negative production cost effect for both oil-importing and oil-exporting countries, but simultaneously a positive income effect for oil-exporting countries. Last, a surge in oil prices increases the wealth of oil-producing countries, thus allowing for increased government spending. In this case, stock returns in an oil-producing country are positively related if private consumption and government spending are complements of each other, and negatively related if private consumption and government spending are substitutes.

3.2 Efficient Markets and Predictability

The previous section illustrates the channels through which oil prices can affect stock prices and provides a background for understanding the mixed relationship found in previous empirical research. The question is, however, if this relationship is fully efficient or if there lies predictability in stock prices based on oil prices, that can be exploited to earn excess returns. The predictability of asset returns has been extensively researched over the years and is still a highly debated topic in financial economics. According to the Efficient Market Hypothesis, stock prices fully reflect all the information available and therefore obtaining abnormal returns is impossible. However, behavioural finance theories among others have challenged the EMH, suggesting that predictable patterns can be identified based on behavioural and psychological factors of stock-price determination.

3.2.1 Efficient Market Hypothesis

The EMH has been primarily derived from concepts presented by Fama (1970) in his famous paper, *"Efficient Capital Markets: A Review of Theory and Empirical Work"*. The EMH describes the market as efficient when the asset prices fully reflect all available relevant information in the market and quickly incorporate news as it becomes known. In other terms, if markets are efficient, then stock prices always trade at their fair market value. This makes it impossible for arbitrage opportunities to arise as there are neither undervalued nor overvalued securities available for trading. Fama's theory is based on the assumptions that all information is free and accessible to all market participants who form rational expectations about future prices and all assess the information in the same fashion.

Fama (1970) suggests three variations of the hypothesis which represent different assumed levels of market efficiency, namely weak, semi-strong and strong form of efficiency. The weak form assumes that the prices of securities reflect all historical information, such as historical sequences of prices and trading volumes, but this might not apply to new information that is not yet publicly available. The semi-strong form assumes that the prices of securities reflect all publicly available information, both historical and new publicly available information, such as dividend announcements and political news. The strong form states that the prices of securities reflect all information. This includes all publicly available information, both historical and current, as well as inside information.

Altogether, according to the EMH, studying movements in past prices and investing based on such observations will not allow investors to earn abnormal returns as historical information is always assumed to be incorporated in the current stock prices under all three forms of efficiency. Additionally, information on past prices cannot help forecast future prices, thus short-term returns should be essentially unpredictable as they follow a random walk.

3.2.2 Random Walk Hypothesis

The EMH is built on the Random Walk Hypothesis, which depicts future price changes as random departures from current prices. The simplest version of the RWH suggests that the best prediction of tomorrow's price is based on today's price plus a random event (i.e., error terms). The dynamics of the stock price can be described as follows:

$$P_t = \mu + P_{t-1} + \varepsilon_t \quad where \qquad \varepsilon_t \sim \ IID\ (0, \sigma^2) \tag{2}$$

where μ is the expected price change, ε_t are random error terms and IID implies that ε_t are independently and identically distributed with mean 0 and variance σ^2. The random error terms are independent, meaning that a shock at time *t-1* does not impact the outcome of the shock at *t*. The fact that they are identically distributed implies that the probabilities for the various outcomes are completely identical, even over time (Campbell et al., 1997).

The rationale behind the RWH is that if information flows properly among the investing public and is fully reflected in the stock prices, then tomorrow's price changes only reflect tomorrow's news and are independent of today's situation. As news is unpredictable, so are price changes. The theoretical argument justifying the RWH and EMH is the assumption that most economic agents are rational and utility maximizing (i.e., seeking high returns) with unbiased expectations of future asset prices. The mispricing that occurs from irrational and uninformed agents should immediately be arbitraged away by rational investors. As a consequence, it is impossible for an investor to consistently beat the market on a risk-adjusted basis since market prices should only react to new information.

3.2.3 Is the stock market efficient?

The EMH has been criticised by many researchers who believe that stock prices can be predicted, at least partially. Empirical results have shown that the EMH is violated in several cases (i.e., Latif et al. (2011); Lamont & Thaler, 2003). The majority of the models used in these empirical studies attempts to show that asset prices are not random walks, which to some extent reveals that asset prices are in fact predictable. Hence, an investor able to predict prices can potentially beat the market.

Challenges to the EMH include a theoretical paradox identified by Grossman and Stiglitz (1980) and several market anomalies. First, Grossman and Stiglitz (1980) show that the EMH entails a paradox, since investors must have an incentive to collect information. If markets are fully efficient, there is no point in active investing because the prices already reflect as much information as one could hope to collect and thus no abnormal returns can be obtained. However, without active investing, who would make the market efficient in the first place? Their point is strengthened by the fact that investors pay large fees to active managers. Thus, either the securities market is inefficient, as active managers can outperform the market, or the market for asset management is inefficient, as investors would pay fees for nothing. It follows that it is logically impossible that all these markets are fully efficient. Second, market anomalies challenge the notion of market efficiency. An example of an anomaly is speculative economic bubbles, which indicate a situation with unsustainable high prices, far from their fundamental value. Bubbles tend to appear in new assets whose fundamental value is hard to assess as there is no prior history of data and are normally detected ex post after a steep fall in prices (Munk, 2019). Therefore, Pedersen (2015) suggests that markets are efficiently inefficient. This term combines the idea that prices reflect all relevant information at all times (i.e., efficiency) with the idea that market prices are significantly influenced by investor irrationality and market frictions (i.e., inefficiency). Although prices are kept in check by intense competition among investors, they are pushed away from their fundamental values because of a variety of demand pressures and institutional frictions. As a result, the market becomes inefficient to an efficient extent. That is, the competition among professional investors makes markets almost efficient, however the market itself remains inefficient enough that they are compensated for their costs and risks. Thus, only a limited amount of capital can be invested with active managers who can beat the market using a few economically motivated investment styles.

The ongoing discussion regarding market efficiency and the shortcomings of the EMH gave rise to a new branch of research, known as behavioural finance. It abandons the idea of market efficiency and argues that behavioural biases induce stock return predictability, as investors are only boundedly rational (Barberis & Thaler, 2002). Examples of behavioural biases left unexplained by the EMH include overconfidence, underreaction, overreaction, representative bias, and information bias. These biases can create predictability and can be exploited to earn excess returns. In the following section, behavioural finance theories and their implications that are relevant for this paper will be presented.

3.3 Behavioural Finance

Behavioural finance is a study of investor market behaviour that focuses on how psychological factors and emotions can affect financial behaviour and market outcomes. Unlike the EMH, behavioural finance assumes that investors are only boundedly rational. Individuals make decisions based on incomplete information and are unable to thoroughly understand and assess its impact. Several social and psychological patterns present in humans (i.e., over- and underreaction) affect investment decisions and consequently stock market efficiency, as they cause prices to deviate from their fundamental values (Shiller, 1999).

3.3.1 Underreaction and Overreaction

The underreaction and overreaction behaviours of investors are primarily derived from misperceptions about information. Underreaction implies that investors do not react enough or in a timely manner to new information, creating a persistent drift in the directional price trend of the stock. On the contrary, overreaction refers to a disproportionate response to new information, which causes the stock to become either overbought or oversold (Barberis et al., 1998). Examples of new information include stock-specific news, firm-specific announcements and information regarding the general economy.

A pioneering study that examines the investor behaviour causing under- and overreaction is conducted by Barberis et al. (1998). They develop a model that attempts to explain under- and overreaction of investors, based on the concept that investors are only boundedly rational. The model is composed of one security and one investor, that reflects the forecast consensus of all investors. The returns of the assets follow a random walk, yet the investor is not aware of this detail and believes that they fluctuate between two states, namely mean-reverting and trending. In the first state, returns shocks are likely to reverse in the following period, such that a positive shock is likely to be followed by a negative one in the next period and then by another positive shock. Conversely, in the second state, shocks are expected to be followed by another shock of the same directionality. After observing the evolution of earnings in a period, the investor updates his or her beliefs on whether the state is mean-reverting or trending. For instance, if good news is followed by better news, then the investor will perceive the state as trending and will invest, thereby creating a momentum effect. This indicates that a series of positive public news – considered as having high strength but low weight – is likely to cause overreaction of stock prices to consistent patterns. This overreaction is due to the fact that investors fail to update their beliefs adequately based on the strength and weight of announcements. The opposite holds for information having low strength but significant weight, such as corporate earnings announcements. In this case, the prediction of stock prices underreacts to this type of information. The research reveals investors' misperceptions about information. In particular, when making forecasts, investors pay too much attention to the strength of the evidence they are presented with and too little attention to its weight.

The observed overreaction and underreaction can be attributed to several behavioural biases. Overreaction of investors to news can be the result of, among others, the herding behaviour, the confirmation bias and representativeness, and overconfidence. Conversely, underreaction can be the result of, among others, the gradual diffusion of information, the anchoring and conservatism bias and the disposition effect. These behavioural biases are covered in the following section.

3.3.1.1 Behavioural Biases causing Overreaction

First, the herding behaviour of investors offers an explanation for overreaction. In the context of stock markets, the herding behaviour refers to the tendency of investors to follow the actions of other investors, rather than individually and independently assessing new information (Cuthbertson & Nitzsche, 2004). Chen (2013) finds that investors are more inclined to herd when information is scarce and difficult to obtain. For instance, domestic investors tend to be more informed about domestic stocks than foreign investors, making foreign investors more likely to herd. Such herding tendency can produce risk beyond what can be justified by information regarding fundamentals and causes prices to deviate (DeLong et al., 1990).

Second, the confirmation bias and representativeness are often offered as potential explanations to overreaction. The confirmation bias is defined as "the seeking or interpreting of evidence in ways that are partial to existing beliefs" (Nickerson, 1998, p.175). Similarly, representativeness implies that individuals tend to solely consider information that is consistent with their existing views and beliefs (Kahneman & Tversky, 1974). In the context of investment decision making, investors consider recent price moves as representative of the future. This attitude can bring them to move capital into investments that have recently made money and conversely out of investments that have recently lost money, causing trends to persist as a result. Pouget et al. (2017) argue that these biases induce overreaction and short-term momentum when traders have either optimistic or pessimistic beliefs. Holding optimistic beliefs is likely to induce biased traders to exaggerate future positive signals, while the opposite holds in case of pessimistic beliefs.

Finally, overconfidence is considered as a possible driver of overreaction. Overconfidence refers to investors' tendency to overestimate their abilities, knowledge and understanding of the equity markets, which can result in unprofitable trading decisions (Plous, 1993). According to Darrat et al. (2007), overconfident investors overestimate the precision of their private signals and trade more than they would if they were fully rational, therefore causing stock prices to deviate from their fundamental values.

3.3.1.2 Behavioural Biases causing Underreaction

First, the underreaction of investors can be explained by the gradual diffusion of information, which refers to the time it takes before private information travels across individuals. The gradual information diffusion hypothesis is theorized by Hong and Stein (1999). They develop a model that features two groups of traders, namely newswatchers and momentum traders. They assume that both types of traders are boundedly rational and react differently to private and public information. Newswatchers make forecasts based on signals that they privately observe and do not condition on past prices. Conversely, momentum traders condition on past price changes. In the simple version of the model, only the newswatchers are present: private information diffuses gradually across the newswatcher's population and prices slowly adjust to new information. The evidence indicates underreaction of prices, which signals predictability, but never overreaction. The situation changes once momentum traders are added in the model. As they condition on past prices, they aim to exploit this underreaction with a simple arbitrage strategy, but they only partially eliminate it. In doing so, they create momentum in prices that culminates in overreaction. Therefore, the existence of underreaction can make it profitable for momentum traders to enter the market and can lead to stock return predictability.

Second, the underreaction of investors can be attributed to the anchoring bias, described as the tendency to rely, or "anchor", one's decisions to the initial piece of information without sufficiently adjusting one's views to new information. This behaviour was initially theorised by Kahneman and Tversky (1974), who conduct experiments in which individuals were first given a random number between zero and a hundred. Then, they were asked to estimate the percentage of African nations in the United Nations. They find that the random number given at the beginning systematically biases the estimations: estimates tend to be higher (lower) for individuals that start with higher (lower) initial numbers. Several research has investigated the anchoring bias in the context of financial markets. For example, George and Hwang (2005) argue that traders tend to use the 52-week high as an anchor against which they evaluate the impact of future news and are reluctant to adjust their bid prices, even in the presence of new positive information.

The anchoring bias is comparable to the conservatism bias, which is defined as the "mental process in which people cling to their prior views at the expense of acknowledging new information" (Pompian, 2006, p.119). This implies that new information is not regarded as important as the beliefs already embedded in investors' minds. This is consistent with

Barberis et al. (1998), who argue that individuals underestimate new information when updating their beliefs. Consequently, prices adjust slow to new information, thereby creating underreaction.

Last, the disposition effect serves as a possible explanation of the underreaction behaviour. The disposition effect was first introduced in finance literature by Shefrin and Statman (1985) and refers to the reluctance of investors to sell assets that have increased in value, while holding onto assets that have lost value. The disposition effect is a widely documented fact in investor behaviour and is examined in several empirical studies. Among others, Frazzini (2006) argues that the disposition effect induces underreaction to news. More precisely, stock prices underreact more to bad news when more current holders are facing a capital loss and underreact more to good news when more current holders are facing a capital gain. Locke and Mann (2000) examine the conduct of professional futures traders and find that all traders hold losers for a longer period than they hold winners, with the least successful traders holding losers the longest.

3.3.2 Underreaction to Information about Oil Prices?

As discussed above, the underreaction of investors to new information is an observed phenomenon in financial markets. As the scope of this study is the effect of oil prices on the stock market, it is of interest to determine whether investors underreact to new information regarding oil prices. This is particularly relevant for investors as underreaction can cause predictability in stock returns and thus can be applied to a trading strategy.

Underreaction to oil price changes is likely to arise as investors might find it challenging to evaluate the impact of such changes on the stock market (Driesprong et al., 2008). The difficulty to assess the impact can be attributed to the fact that oil prices impact stock returns through a variety of channels. The combination of effects within a specific channel or across different channels can often be contrasting and thus generate ambiguous responses from the stock markets. Generally speaking, investors are concerned in understanding solely the market they specialize in and are unable to devote attention to processing potential valuable news from other markets on time (Merton, 1987). Thus, investors pay more attention to news regarding oil prices when they are trading assets that have a direct link to oil prices (Pollet, 2011). This explains why in oil-dependent markets the impact of oil price changes is more accurately assessed and therefore underreaction is less pronounced (Fan & Jahan-Parvar, 2011). On the contrary, assessing the effect of oil price changes is more difficult in non-oil related markets, thus increasing the underreaction.

The delayed reaction of investors challenges the notion of market efficiency. If markets were efficient, one would expect that information about oil price changes would be immediately incorporated in stock prices, especially since traders in every market have access to oil prices on a real time basis. However, due to investors' limited information-processing capacity, it can take time before information about oil price changes is fully reflected in the stock markets. Altogether, it is expected that there is a delayed reaction of investors to information about oil prices.

4. Hypothesis Development

As discussed in the previous sections, several empirical studies are conducted with the purpose of investigating the relationship between oil prices and stock prices. However, there is a lack of consensus regarding the directionality of the relationship. With these premises, we believe it is of interest to research this relation. In particular, we aim to investigate whether changes in oil prices have an effect on stock market returns. This leads to the following hypothesis:

Hypothesis 1: Oil prices do not have an effect on stock market returns.

Which can be formalized into a testable hypothesis as:

H_0: $\beta^{oil}= 0$

H_A: $\beta^{oil} \neq 0$

where β^{oil} is the coefficient of the oil factor that is estimated using a rolling-window regression, as defined below:

$$r_t^{stock} = \alpha + \beta^{oil} * r_t^{oil} + \beta^{IPI} * r_t^{IPI} + u_t \quad (3)$$

The chosen regression model allows us to capture the time-varying relationship between oil prices and stock markets. To test this first hypothesis, we assess whether the coefficient of the oil factor is significantly different from zero. If the oil coefficient is significantly different from zero, we reject the null hypothesis that oil prices do not have an effect on stock market returns, in favour of the alternative. On the contrary, if the coefficient is not significantly different from zero, we fail to reject the null hypothesis of no oil effect. We expect to find significant linkages, at least to some extent, between oil prices and stock returns based on previous research.

After establishing whether a relationship between oil prices and stock returns exists and how this relationship evolves over time, we are interested in determining whether the oil effect strengthens when investors' underreaction is taken into consideration. As argued by Hong and Stein (1999), it takes time before information about oil price changes is fully reflected in the stock markets. To examine if investors react with a delay to oil prices the following hypothesis is formulated:

Hypothesis 2: Investors do not react with a delay to oil prices.

We test whether there is underreaction by introducing lags of several trading days between monthly oil returns and stock returns and then by comparing the explanatory power of the regressions. Each regression is estimated according to the following equation:

$$r_t^{stock} = \alpha + \beta * r_{t-j}^{oil} + +u_t \quad (4)$$

where *j* represents the number of lagged days. Assuming there is a delayed reaction, introducing a lag between the monthly oil price data and stock returns should improve the explanatory power of the regression. This explanatory power is expected to decrease once the majority of the information is incorporated in stock prices. In line with the theoretical review and the findings of Driesprong et al. (2008) and Fan and Jahan-Parvar (2011), we expect to find a delayed reaction of investors to oil price changes. The findings of the speed of investor reaction, i.e., the lag with the highest explanatory power, can be applied in a trading strategy to attempt to predict stock returns and thus provide practical implications for investors.

5. Data

This section outlines the data employed in the empirical research. As the focus of the analysis is to model the relationship between oil prices and stock market, two primary series of data are used, namely WTI Crude Oil Futures and stock market indices. Further, in order to better capture the relation, the industrial production index is included as a control variable. Each section provides a brief explanation of the data, combined with plots to illustrate the dynamics of the series and the summary statistics.

5.1 Oil Price Data

West Texas Intermediate (WTI) and Brent are the two primary references for crude oil prices globally. WTI is adopted primarily in North America as a benchmark, whereas Brent is the benchmark in Europe, Africa and Middle East. Asian countries adopt a mixture of WTI and Brent prices to value their crude oil. As illustrated in Figure 1, the dynamics of WTI and Brent prices move very closely together. The graph plots the spot prices for both benchmarks, expressed in terms of USD per barrel. The two oil benchmarks move in tandem, where WTI historically traded at a premium over Brent with a price difference fluctuating around three USD per barrel (Caporin et al., 2019). This is confirmed by the high correlation between Brent and WTI of 0.9786 observed over the period under analysis. The proximity of prices and their convergence to their long-term pattern is associated with a globalised market. In a globalised and thus integrated market, supply and demand shocks affecting the price of petroleum in one region immediately shift to other markets. The proximity in prices can in addition be attributed to the low transportation costs in the oil market. Unlike other energy commodities, petroleum can be quickly and cheaply transferred across locations due to the efficient pipelines and infrastructures. As a consequence, the prices of crude oil of the same quality but different locations move together. This synchronous relationship was broken in the beginning of 2011, when WTI began to trade at a large discount compared to Brent, as can be seen in Figure 1. Complications in the U.S. oil transportation infrastructure led to difficulties in moving crude oil from a storage location in Cushing, Oklahoma, to the refineries in the Gulf Coast. This resulted in an oil surplus that depreciated WTI compared to Brent (Büyükşahin & Robe, 2011). In 2014, these infrastructure problems in the WTI market started to decrease, leading to a recoupling of the two crude oil benchmarks. In recent years, however, WTI trades at a small discount compared to Brent.

In this analysis, only WTI is used as a proxy for oil prices, as it has the largest data availability. This is due to the fact that WTI started being traded earlier than Brent. As the prices of both Brent and WTI are closely correlated and only differ by a small amount, this choice should not make a notable difference in the results.

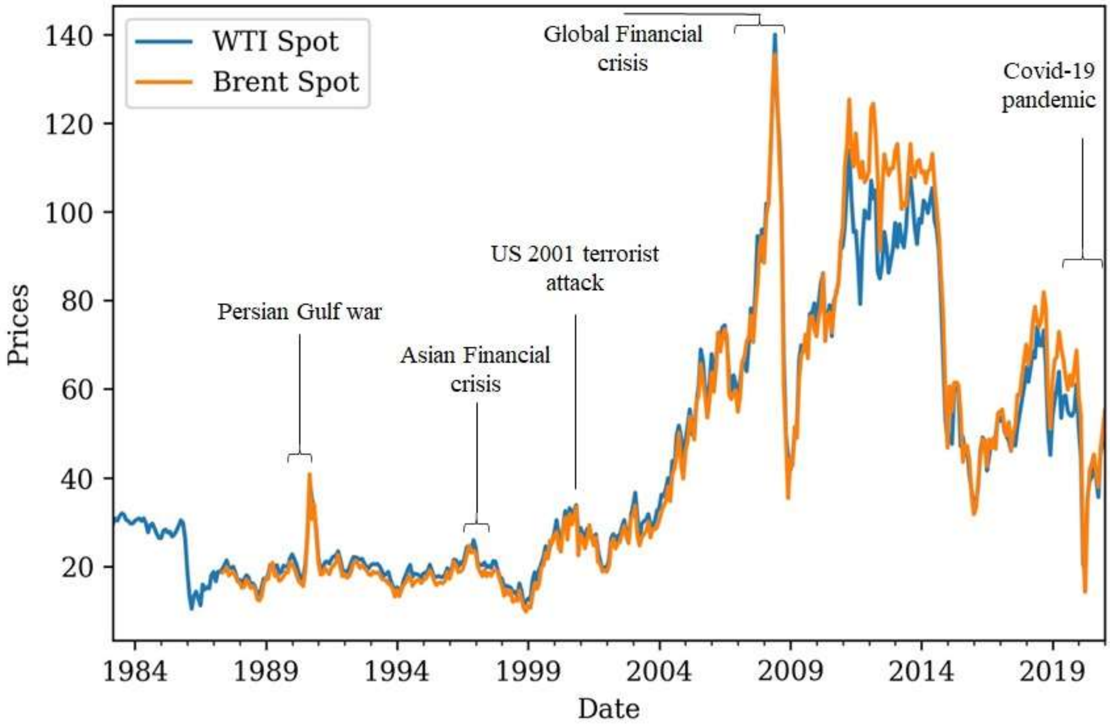

Figure 1: Brent and WTI Spot Oil Prices. Source: Bloomberg

A typical feature of oil markets is the high level of oil price volatility and the fact that their movements are greatly affected by exogenous events. Figure 1 illustrates the dynamics of both WTI and Brent spot price series, where several major oil price shocks can be identified. For example, the Persian Gulf War, that started in 1990/1991 in response to Iraq's invasion of Kuwait, caused disruptions in the oil supply as the affected area accounted for nearly 9% of the global oil production. The prices declined to their pre-shock levels relatively quickly, as the excess capacity of Saudi Arabia helped restore the oil production levels (Hamilton, 2013). In 1997/1998, the Asian financial crisis led to an oil price drop of around 50%. The financial crisis started in Thailand after its currency, the Thai Baht, plunged in value and soon spread to other Southeast Asian countries including China, the second largest oil-consumer at that time after the United States. The crisis caused a drop in demand for oil, which resulted in a decrease in oil prices from $20 to $13 per barrel (Wang et al., 2013). Prices began to recover in 1999 and started stabilizing in 2000 when OPEC[2] overshot the oil production and cut quotas. In September 2001, the terrorist attack on the World Trade Center in New York caused panic and fear of air travel. Consequently, demand for oil decreased and crude oil prices fell drastically: the spot prices for WTI were down by 35% in November 2001. The OPEC decision to cut oil production quotas in 2002 pushed however the prices up again (Sadorsky, 2012). The global financial crisis of 2008/2009 resulted in substantial instability in the financial markets, which in turn affected the oil market. A stagnant oil supply and lower demand caused the oil prices to increase to $150 per barrel in July 2008 and fall below $40 at the end of 2008. In 2015, the oil prices fell to its lowest level since the global financial crisis. This major drop has been caused by a lower demand for oil in Europe and China, coupled with a booming U.S. oil production as well as a steady supply of oil from OPEC (Camp et al., 2020). Crude oil prices registered another dramatic decrease in 2020, where petroleum prices plummeted below zero. This can be primarily attributed to the Covid-19 pandemic, which led to a global economic slowdown and to significant shocks in oil supply and demand.

When analysing the relationship between oil prices and the stock market, either crude oil spot prices or futures prices can be used. The spot price is the current price in the marketplace at which an asset can be bought or sold for immediate delivery, whereas the futures price is the price for which the asset can be bought or sold at a certain future time. Despite some differences, WTI spot and futures prices generally move closely together, as can be seen in Figure 2. In this paper, futures prices are used instead of spot prices, as it can be argued that futures prices are a better indicator of crude oil price than spot prices. This is mainly due to the fact that the futures markets dominate the spot markets at most points in time due to high volumes of speculation in the derivatives market. Namely, speculators can either invest in the spot market with the real commodity by buying and selling an actual barrel of oil, or they can invest in the futures market, which is a more common activity (Zhang & Wang, 2013). Thereby, the low transaction costs and wide use of short-selling mechanisms in the futures market help the futures prices react quicker to new information which makes them more efficient (Alquist & Kilian, 2010).

[2] OPEC stands for Organization of the Petroleum Exporting Countries. OPEC was created in September 1960, with the purpose of coordinating and unifying petroleum policies among Member Countries, securing stable prices and an efficient and regular supply of petroleum to consuming countries. Currently, OPEC is composed of the following countries: Algeria, Angola, Congo, Equatorial Guinea, Gabon, Iran, Iraq, Kuwait, Libya, Nigeria, Saudi Arabia, United Arab Emirates and Venezuela (OPEC, n.d.).

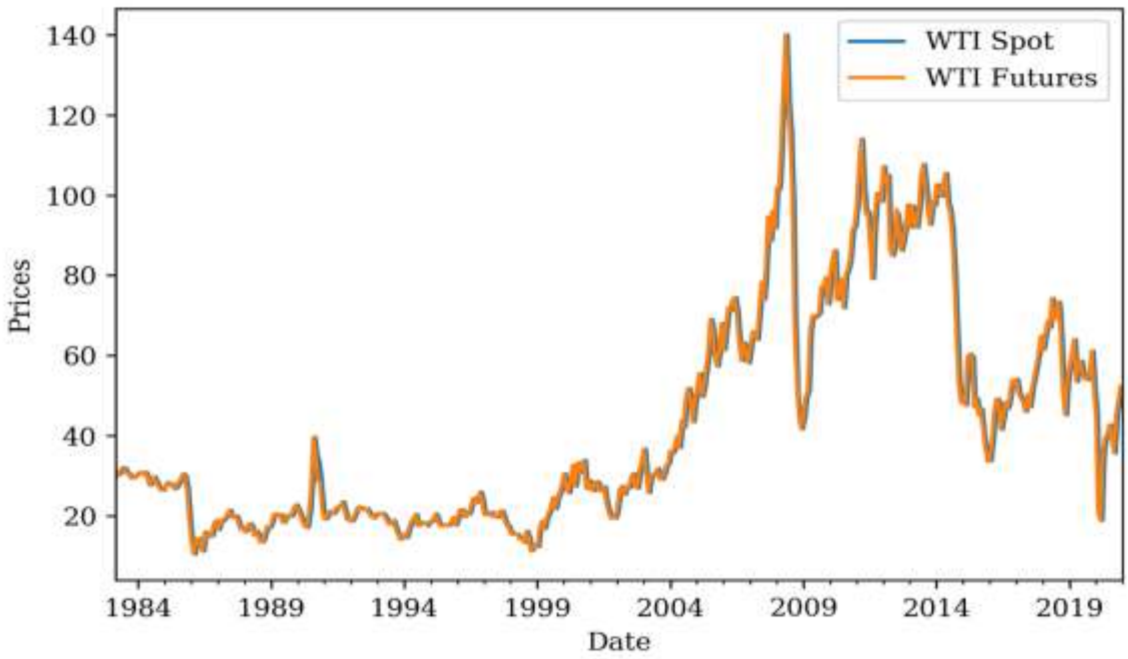

Figure 2: WTI Spot and Futures Prices

Ideally, the analysis would start in 1973, as the oil crisis in that year caused the prices to fluctuate. However, WTI futures contracts only started being traded on the Chicago Board of Trade in 1983. Therefore, the period under analysis spans from March 1983 to December 2020. Monthly data is used as this tends to be less noisy than daily data (Brooks, 2008). Continuously compounded monthly returns are calculated by taking the natural logarithm of the stock price at time *t* divided by the stock price at *t-1*, as shown in the following equation:

$$R_{i,t} = \ln\left(\frac{P_{i,t}}{P_{i,t-1}}\right) \quad (5)$$

where *ln* is the natural logarithm, $P_{i,t}$ is the stock price of firm *i* at time *t* and $P_{i,t-1}$ is the stock price of firm *i* at time *t-1*. The logarithmic transformation makes returns time-additive, meaning that the logarithmic return over a period of *t* to *T* is the sum of all logarithmic returns between the *t* and *T*. Additionally, transforming to logarithmic returns normalizes the variables, implying that all returns are in a comparable metric. Lastly, logarithmic returns eliminate the non-stationary properties of the returns (Brooks, 2008).

Table 1 contains information about the oil price series used and Table 2 presents the descriptive statistics of the WTI futures series. Futures prices have been retrieved as end-of-month data. As it is often the case for logarithmic returns, the mean is close to zero, as, by construction, logarithmic returns revolved around zero. When analysing descriptive statistics, kurtosis and skewness are two important metrics to take into account. Kurtosis measures whether the data is heavy- or light-tailed relative to the kurtosis value of a normal distribution, which is equal to three. The series presents an excess positive kurtosis, which implies that the returns are not normally distributed. In such a scenario, we expect broader fluctuations in the data, resulting in greater probability of extreme outlier values, both positive and negative. Skewness measures the symmetry in a dataset, where perfectly symmetrical data has a skewness of zero. WTI futures prices show moderate negative skewness of -0.5724, which implies that the distribution of returns is not symmetrical (Brooks, 2008).

Table 1: Oil Price Series

Short name	Full name	Source	Starting Date
WTI Future	NYMEX West Texas Intermediate continuous set.	Bloomberg	March 1983

Table 2: Descriptive Statistics of Oil Price Futures, expressed as logarithmic returns

	Obs.	Mean	SD.	Min	25%	50%	75%	Max	Kurt.	Skew.
WTI	453	0.12%	10.60%	-0.782	-0.053	0.0088	0.0608	0.6333	9.7837	-0.572

5.2 Stock Market Data

The dataset comprises stock market data of six primarily oil-exporting countries (Canada, Russia, Norway, Brazil, Saudi Arabia and Mexico) and four primarily oil-importing countries (United States, China, Jordan and Hong Kong). This diverse set of countries is chosen following three criteria. First, the sample includes a mixture of oil-importers and oil-exporters. Second, it includes both developed and emerging stock markets. Third, the countries belong to different geographies. Altogether, these criteria allow for an in-depth analysis of the relationship between equity markets and oil prices as they capture possible heterogeneous responses depending on the characteristics of the country and across different geographies. Similarly to the oil price data series, stock market data are collected as end-of-month observations for the period March 1983-December 2020. Table 3 provides information about the stock indices used in the analysis. Stock prices are expressed in the country's local currency and differences in sample periods are due to data availability.

Table 3: Stock Market Indices

Country	Index	Currency	Source	Starting Date
United States	S&P 500	USD	Datastream	March 1983
Mexico	IPC Mexico (BOLSA)	MXN	Datastream	January 1988
Hong Kong	Hang Seng	HKD	Datastream	December 1992
Canada	S&P/TSX Composite	CAD	Datastream	March 1983
China	SSEA	CNY	Datastream	January 1991
Norway	Oslo OBX	NOK	Datastream	January 1987
Jordan	Amman Stock Exchange	JOD	Datastream	December 1994
Saudi Arabia	SASEIDX	SAR	Bloomberg	January 1994
Brazil	Bovespa	BRL	Datastream	February 1991
Russia	MOEX Russia	RUR	Bloomberg	September 1997

Figure 3 illustrates the historical dynamics of the countries' stock market indices. Generally, all countries exhibit positive trends throughout the time period, with emerging markets (such as China, Mexico, Brazil and Russia) recording the largest increase in index value. Emerging countries have experienced rapid growth in recent years, in terms of market capitalization, productivity and GDP per capita. In addition to the large growth potential, emerging markets also present greater risks, such as political instability, currency fluctuations and higher cost to invest. These risks result in higher stock market volatility, as equity indices more frequently experience either upturns or downturns. Figure 3 also shows that several exogenous circumstances affect all stock markets to some extent. For instance, all stock indices sustained a sudden and large decline in value in 2008/2009, following the financial turmoil that started in the United States. Similarly, between February and April 2020, global equity markets entered a bearing position after the start of the Covid-19 pandemic.

Table 4 displays the descriptive statistics of the stock market returns. All countries exhibit a positive mean return, with Brazil, Mexico and Russia recording the largest average (3.91%, 1.46% and 1.25% respectively). In general, equity returns of emerging markets show higher volatilities compared to established markets. More precisely, Brazil, China and Russia report a high standard deviation of 13.74%, 11.65%, and 10.79% respectively. Several return series show relatively high kurtosis, in particular China, Hong Kong, Canada, Brazil and Russia. As mentioned previously, excess positive kurtosis signals the presence of extreme outlier values, both positive and negative. Except Saudi Arabia, all countries present a kurtosis higher than the normal distribution. As far as the skewness is concerned, all nations present negative values, except for Jordan, Brazil and China. The overall excess kurtosis and negative skewness in the stock return series is expected and in line with research that has shown that the stock market does not move parallel with the normal distribution. These so-called stylized facts are statistical properties that appear to be present in many empirical assets returns across time and markets (Brooks, 2008). The correlation between the index series tends to be positive and low, as can be seen from Table A1 in Appendix A.

Table 4: Descriptive Statistics of the Stock Return Series, expressed as logarithmic returns

	United States	Mexico	Hong Kong	Canada	China	Norway	Jordan	Saudi Arabia	Brazil	Russia
Obs.	453	395	453	453	359	407	385	323	358	279
Mean	0.71%	1.46%	0.73%	0.46%	0.92%	0.55%	0.35%	0.50%	3.91%	1.25%
SD.	4.38%	7.33%	7.63%	4.27%	11.65%	6.60%	4.63%	6.57%	13.74%	10.79%
Min	-0.2454	-0.3498	-0.5656	-0.2566	-0.3856	-0.3444	-0.2121	-0.2978	-0.5034	-0.5826
25%	-0.0165	-0.0237	-0.0294	-0.0151	-0.0471	-0.0253	-0.0228	-0.0263	-0.0315	-0.0333
50%	0.0113	0.0141	0.0123	0.0085	0.0080	0.0137	-0.0008	0.0099	0.0178	0.0167
75%	0.0346	0.0539	0.0488	0.0303	0.0556	0.0459	0.0262	0.0403	0.0821	0.0603
Max	0.1238	0.3623	0.2645	0.1119	0.9771	0.1696	0.2115	0.1790	0.6931	0.4255
Kurt.	3.3850	3.1833	8.1125	6.3828	15.5956	4.2465	3.3611	2.3744	6.1443	5.9132
Skew.	-0.9757	-0.2259	-1.1697	-1.4053	1.8774	-1.3306	0.3281	-0.7252	1.4953	-0.8568

Figure 3: Stock Market Indices

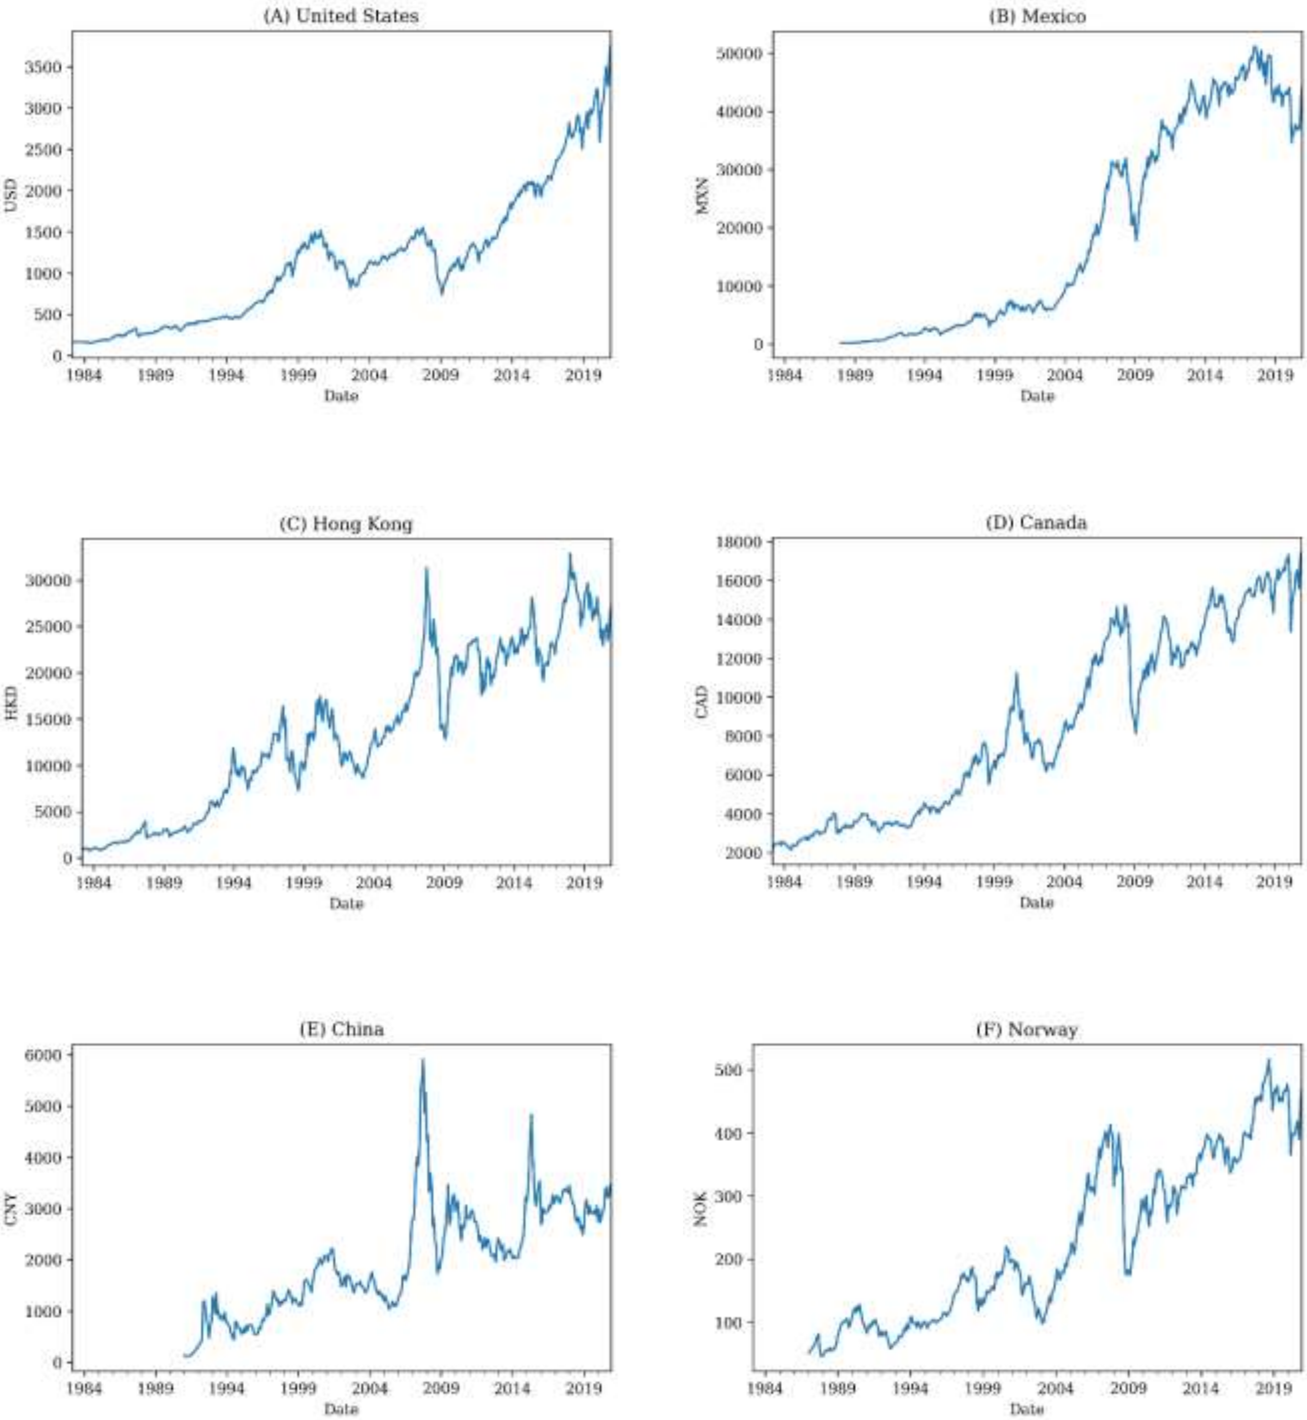

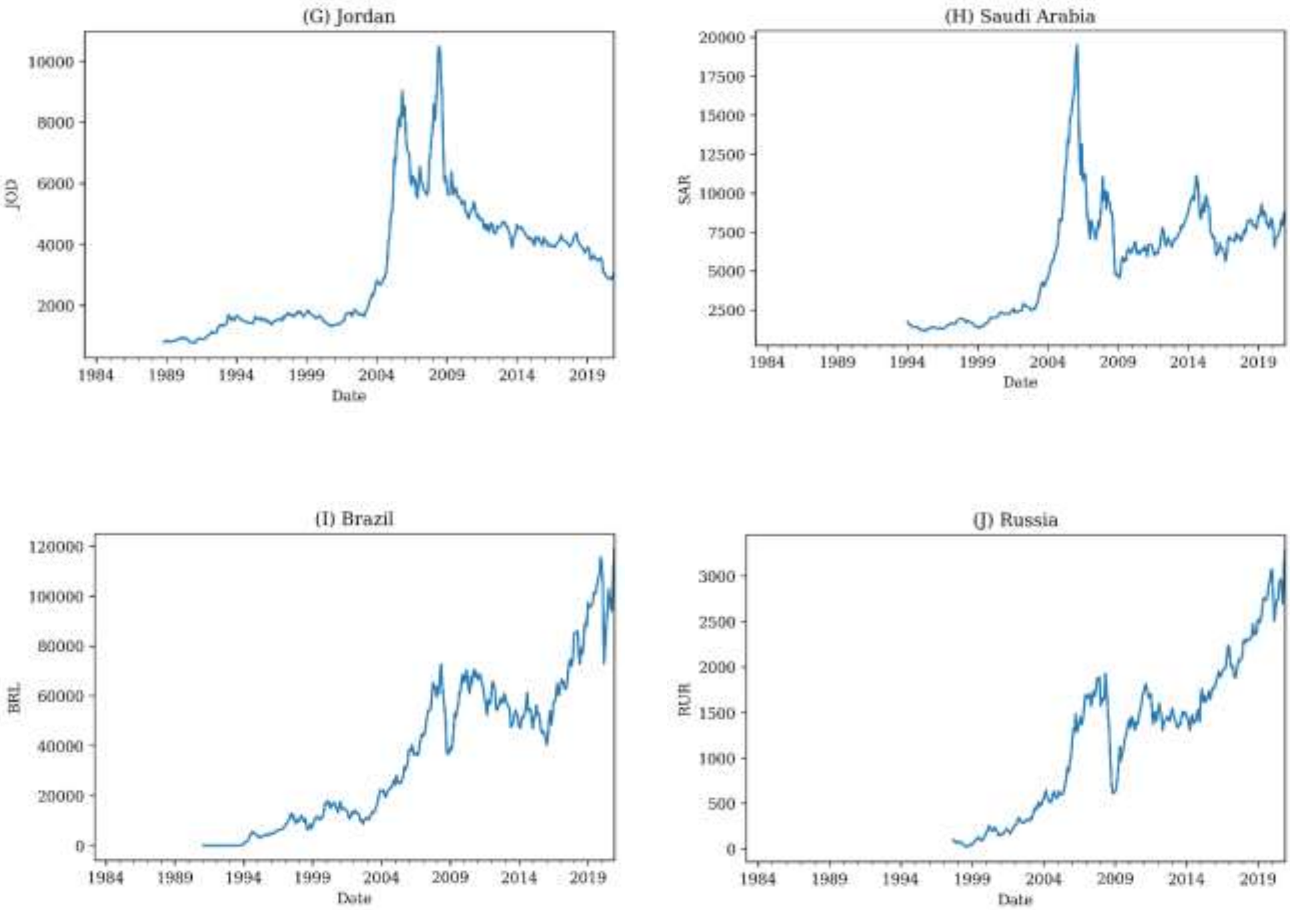
(G) Jordan
JOD
10000
8000
6000
4000
2000
1984
1989
1994
1999
2004
2009
2014
2019
Date
(H) Saudi Arabia
SAR
20000
17500
15000
12500
10000
7500
5000
2500
1984
1989
1994
1999
2004
2009
2014
2019
Date
(I) Brazil
BRL
120000
100000
80000
60000
40000
20000
0
1984
1989
1994
1999
2004
2009
2014
2019
Date
(J) Russia
RUR
3000
2500
2000
1500
1000
500
0
1984
1989
1994
1999
2004
2009
2014
2019
Date

5.3 Industrial Production Data

The industrial production index (IPI) measures the level of production in the manufacturing, mining, electric and gas industries in the United States on a monthly basis (Federal Reserve Bank of St.Louis, n.d.). Changes in the IPI are considered as a relevant indicator that reflects similar changes in the overall economic activity. In particular, the IPI is a leading indicator of GDP growth and economic performance due to its sensitivity to consumer demand and interest rates. IPI is also used by central banks to measure inflation, as high levels of industrial production can lead to uncontrolled levels of consumption and thus rapid inflation (Shapiro, 1989). An increase in industrial production is expected to raise the future cash flows and thus the profitability of firms. Therefore, the relationship between industrial production and stock returns is expected to be positive, as argued by Balvers et al. (1990) and Schwert (1990). The IPI is included in our research as a control variable as it is a good proxy for the overall economic activity. Including a control variable is common practice in econometrics and regression analysis as it usually improves the fit of the model and the reliability of the estimated coefficients (Wooldridge, 2008). In this paper, the use of the IPI as a control variable enables us to test the impact of oil price changes on stock returns, while controlling for the effect of IPI on stock returns.

In this analysis, the IPI of the United States has been used as a proxy for all the countries in the sample, as the United States has one of the largest and most influential economies and thus its industrial production growth is highly correlated with the world's industrial production growth. The index level data are expressed as the volume of real output and activity in each month, relative to a base year. Figure 4 plots the dynamics of the change in IPI throughout the period under analysis. The plot shows an overall upward trend, which is in line with the boost that the industrial sector in the United States has experienced over the years. Moreover, the plot illustrates significant drops in industrial production during the financial crisis of 2008/2009 and during the recent Covid-19 crisis, as periods of recessions are characterized by a decrease in the industrial production level.

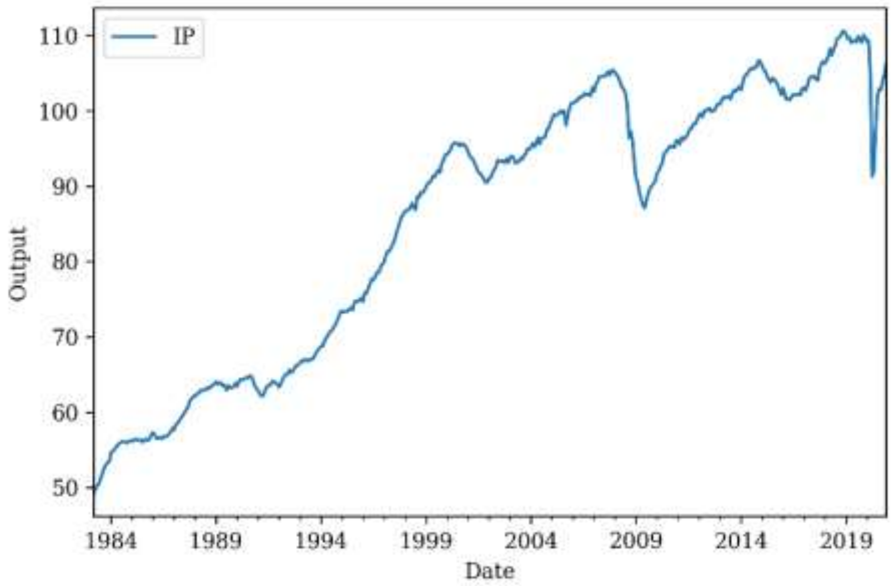

Figure 4: Industrial Production Index

Similarly to oil and stock market data, the original price series has been transformed into logarithmic returns. Table 5 contains information of the IPI data and Table 6 reports the summary statistics of the IPI returns. As it is often the case for logarithmic returns, the mean revolves around zero. The excess kurtosis value shows that the series is prone to extreme outcomes, while the negative skewness indicates that the returns are not normally distributed.

Table 5: Industrial Production Series

Short name	Full name	Source	Starting Date	Currency
IP Index	US Industrial Production Index	Bloomberg	March 1983	USD

Table 6: Descriptive Statistics of Industrial Production Index, expressed as logarithmic returns

	Obs.	Mean	SD.	Min	25%	50%	75%	Max	Kurt.	Skew.
IPI	453	0.17%	0.98%	-0.136	-0.002	0.0022	0.0057	0.0605	90.945	-6.103

6. Methodology

This section presents the methodology used to address the hypotheses of this paper. Section 6.1 presents the concept of stationarity and the statistical tests used to test this property. Section 6.2 explains the Ordinary Least Squares (OLS) regression model with its underlying assumptions. Section 6.3 discusses the different approaches used to test for the two hypotheses, where for the first hypothesis a rolling OLS regression will be estimated, and for the second hypothesis a static OLS regression will be estimated.

6.1. Stationarity

6.1.1 Definition

Stationarity is a core concept in econometrics and time series analysis as various models and statistical tools rely on it. The stationarity, or non-stationarity, of a time series greatly affects its behaviour and properties. A series follows a strictly stationary process if its statistical properties, which are mean, variance and autocovariance, remain constant over time. For a time series to be employed in regression analysis, weak (or covariance) stationarity is sufficient. Stationarity is a prerequisite in time series analysis as it enhances the reliability of the results. Working with non-stationary data is likely to lead to spurious regressions[3]. These generally return high explanatory power (R^2) indicating that the series are strongly correlated, even though a real relationship between the series does not exist (Brooks, 2008).

6.1.2 Testing for Stationarity

There are several ways of testing for stationarity. The simplest approach consists in plotting the data and determining visually whether the series presents some characteristics of non-stationary data, such as seasonality or trend. A second approach involves looking at the autocorrelation function (ACF), which measures the relationship between a variable's current value and its past (or lagged) values (Brooks, 2008). When plotting the ACF, the values tend to decrease to zero quickly for stationary processes, while the opposite holds for non-stationary data. A third

[3] A spurious regression provides misleading evidence of a linear relationship between two non-stationary and independent time series. The regression wrongly estimates the parameters, providing significant t-values and high R-squared (Brooks, 2008).

method is to use statistical procedures. There exist two types of statistical tests, the unit root and stationarity tests, and they differ in the way the underlying hypotheses are specified. Unit root tests examine whether a time series is non-stationary (i.e., possesses a unit root), whereas stationarity tests examine whether a time series is stationary. In this paper, two approaches have been adopted: the Augmented Dickey-Fuller (ADF) test and the Kwiatkowski-Phillips-Schmidt-Shin (KPSS) test.

The ADF test, which is the most common unit root approach, can be expressed by the following formula:

$$\Delta y_t = \mu + \delta y_{t-1} + \sum_{i=1}^{p} \alpha_i \Delta y_{t-1} + u_t \tag{6}$$

where Δy_t is the first difference of the series at time t-1, y_{t-1} is the first lag of the time series (y) and δ is the coefficient of the first lag of y. The null hypothesis is that the time series has a unit root (H_0: δ=0), whereas the alternative hypothesis is that the series is stationary (H_A: δ<0). At a 95% confidence interval, the null hypothesis is rejected in favour of the alternative if the p-value is less than 0.05, thereby inferring that the series does not have a unit root and is stationary. On the contrary, if the p-value is greater than (or equal to) 0.05, the null hypothesis fails to be rejected and the series is non-stationary.

Conversely to the ADF test, the KPSS test is used to check the stationarity of a time series. The test is derived from the following linear equation:

$$x_t = r_t + \beta_t + \varepsilon_t \tag{7}$$

where r_t is a random walk process, β_t is a deterministic trend and ε_t is the error term. The random walk process can be expressed as:

$$r_t = r_{t-1} + u_t \tag{8}$$

where u_t ~ (0,ϑ^2) and are IID. The null hypothesis (H_0: ϑ^2 = 0, assumes stationarity around a mean or linear trend, while the alternative (H_A: ϑ^2> 0) indicates the presence of a unit root. If the t-statistic is greater than 0.463 (which is the critical value for the 95% confidence interval), the null hypothesis of stationarity is rejected in favour of the alternative. On the contrary, if the t-statistic is lower than (or equal to) 0.146, the test fails to reject the null hypothesis and the series is stationary.

In empirical analysis, the KPSS test is often used to complement the ADF test. The joint results provide a better understanding of the statistical properties of the series. In case the series under analysis possess a unit root, some approaches can be taken in order to make them stationary. Logarithmic transformation is a commonly used solution, as it allows to remove any trend and stabilize the variance for a series.

6.2 Regression Analysis

Regression analysis is widely used in empirical research with the purpose of estimating the relationship between a dependent variable (*y*) and one or more independent variables (*x*). Linear regressions attempt to model the linear relation between *x* and *y* by fitting a straight line to the data. *Y* can be expressed as a linear function of *x*, as illustrated in the following formula:

$$y_t = \alpha + \beta * x_t + u_t \tag{9}$$

where α is the intercept, β is the slope of the line and u_t is the error term. The β coefficient measures the variation in the dependent variable when the independent one has a unitary variation. The error term u_t is added to capture stochastic deviations in *y* that the explanatory variable fails to explain. Depending on the number of independent variables included, the linear regression can be defined as either a simple (if there is only one *x*) or a multiple (in case of two or more x) regression model (Brooks, 2008).

Ordinary Least Squares (OLS) is a method used to estimate the relation between dependent and independent variables. OLS determines the parameters of the regression by following the principle of least squares, which implies minimizing the sum of squared errors, defined as the difference between the estimated $\hat{y}$ and observed y (Brooks, 2008). The sum of squared differences (or errors) is generally referred to as the residual sum of squares (RSS). The OLS approach identifies the estimated $\hat{\alpha}$ and $\hat{\beta}$ that minimize RSS, as shown in the following equation:

$$RSS = \sum_{t=1}^{T} \left(y_t - \hat{\alpha} - \hat{\beta} * x_t\right)^2 \tag{10}$$

Generally speaking, a model fits the data well if the variation between the observed and fitted values is small and unbiased. Hence, a low RSS is often preferred.

To test for the first hypothesis, a rolling-window regression is adopted, as we are modelling the contemporaneous and time-varying relationship between dependent and independent variables at time t. To test for the second hypothesis, a static OLS regression is estimated to test for investors' underreaction.

6.2.1 OLS Assumptions

To ensure that the output of the regression is accurate and the estimates can be used to test the hypotheses, the OLS regression needs to satisfy the following assumptions (Brooks, 2008):

$$E\ (u_t) = 0 \tag{11}$$

$$Var\ (u_t) = \vartheta^2 < \infty \tag{12}$$

$$Cov\ \left(u_i, u_j\right) = 0 \tag{13}$$

$$Cov\ (u_i, x_t) = 0 \tag{14}$$

The first assumption indicates that the average value of the error term is zero. The error term accounts for variation in the dependent variables that the independent factors fail to explain. For the model to be unbiased, the average value of the error should be zero. A value different from zero implies that the model statistically unpredicts the fitted values. The inclusion of a constant term (α) in the regression eliminates the risk of violating this first assumption, by forcing the mean of the residuals to equal zero. In particular, α corresponds to the portion of the dependent variable that cannot be explained by the independent variables.

The second assumption is that the variance of the errors must be constant over all observations of the independent variables. This characteristic is known as homoskedasticity. If the errors do not have a constant variance, thus they are heteroskedastic, they can lead to wrong inferences and biased standard errors. One way of overcoming the issue of heteroskedasticity is to transform the variables into natural logarithms or reduce variables by some other measures of 'size', which has the effect of re-scaling data to 'pull in' extreme observations. To account for heteroskedasticity in the error terms, Newey-West covariance matrix estimator is used to convert the standard errors into Heteroskedasticity and Autocorrelation Corrected (HAC) standard errors.

The third assumption postulates that the covariance between the error terms over time is zero, that is, the errors are uncorrelated with one another implying no autocorrelation. Autocorrelation measures the relationship between the residual current values and its past values. Working with autocorrelated residuals will result in biased standard errors. Similarly to the case of heteroskedasticity, the Newey-West estimator is used to convert the standard errors into HAC standard errors to overcome autocorrelation.

The fourth assumption indicates that all independent variables are uncorrelated with the error term (exogeneity). Correlation between an independent variable (endogeneity) and the errors implies that the former can be used to predict the latter, thus violating the notion that errors are stochastic and unpredictable. Endogeneity arises when the dependent variable is a predictor of the independent variable and not only a response to it. As argued by Ciner (2001), the influence between stock returns and oil prices is mutual, as stock index returns have an impact on crude oil futures prices. When endogeneity between the dependent and independent variable is a natural result of their theoretical relationship, it becomes increasingly difficult to draw conclusions about the direction of causality and the regression results become difficult to interpret. For example, Barsky and Kilian (2001) argue that oil price shocks happen in connection with political instability and conflicts, which simultaneously affects stock markets. To attempt controlling for exogenous events, the industrial production factor is included in the regression model. The endogeneity problem occurs in many areas in economic research due to the nature of economics, as economic variables are highly interrelated with each other. However, fully accounting for the endogeneity problem is rather complicated in non-experimental settings when macroeconomic variables are involved. Therefore, there is some likelihood that the results of this research are biased and thus should be interpreted with caution.

6.3 Empirical Analysis Hypotheses

6.3.1 Oil Effect

In the first section of the analysis, we investigate the contemporaneous relationship between oil prices and stock market. As outlined in Section 2.2, empirical evidence shows that this relationship is time-varying. A common assumption of time series analysis is that the model's parameters are time-invariant. However, as the economic environment often changes considerably, it may be more reasonable to assume in certain settings that the parameters evolve over time. The rolling approach is often used to model the changing relationship between a dependent variable and one (or more) explanatory variables and to assess the constancy of the parameters over time. As this part of the analysis focuses on investigating the time-varying relationship between stock prices and oil prices contemporaneously, a rolling-window regression analysis is implemented. This way it is possible to obtain an overview of the evolution of the beta coefficient throughout the time period.

In the rolling-window regression, two explanatory variables are included, namely the oil price and the industrial production factor. The oil factor is incorporated in the regression analysis to test for the existence of an oil effect, while the IPI factor serves as a control variable. As oil price shocks are likely to be driven by exogenous events, such as political instability or economic recessions, the IPI is included in the analysis in order to attempt controlling for these events. Including a control variable usually improves the fit of the model, expressed by the R^2 or adjusted R^2. These statistics indicate the percentage of the variance in the dependent variable that can be collectively explained by the independent variables. The inclusion of a control variable increases the reliability of the estimated coefficients, as the variability due to unexplained variables should be excluded (Wooldridge, 2008). However, including extra control variables does not necessarily lead to a better explained model as it can cause the model to overfit. Therefore, in order to limit the risk of overfitting, only one control variable is added to the model.

The rolling-window method performs multiple regressions, with subsamples of the original sample. To obtain these subsamples, a window of a certain size is defined and kept constant. A regression is performed on the observations contained in the window, then the window is moved one observation forward and the process is repeated throughout the entire sample. Thus, the parameters of interest are estimated across different sampling periods with the same window size. For a window of width $n < T$, the rolling regression model can be expressed as:

$$y_t(n) = X_t(n) * \beta_t(n) + u_t(n) \; with \; t = n, \ldots, T \tag{15}$$

where $y_t(n)$ is a vector of observations on the dependent variable, $X_t(n)$ is a matrix of explanatory variables, $\beta_t(n)$ is a vector of the estimated parameters and $u_t(n)$ is a vector of error terms. To test the first hypothesis, a rolling-window OLS regression is estimated for each stock index, according to the following equation:

$$r_t^{stock} = \alpha + \beta^{oil} * r_t^{oil} + \beta^{IPI} * r_t^{IPI} + u_t \quad (16)$$

where r_t^{stock} is the logarithmic return of the country index at time *t*, α is the constant term, β^{oil} is the estimated coefficient of the oil factor, β^{IPI} is the estimated coefficient of the IPI factor, r_t^{oil} is the log logarithmic in oil futures prices at time *t*, r_t^{IPI} is the logarithmic change in the IPI at time *t* and u_t is the standard error. Each OLS regression yields estimates for the intercept of the regression (α) and coefficients (β) of both the oil and industrial production factors. Further, the estimated coefficients have been used to obtain expected stock return values, in accordance with the following formula:

$$\widehat{y_t} = \widehat{\alpha_t} + \widehat{\beta_t} * r_t^{oil} + \widehat{\beta_t} * r_t^{IPI} \quad (17)$$

The performance of the rolling regression is sensitive to the choice of window size (Inoue & Rossi, 2012). In general, shorter rolling window sizes are preferred for data collected in short intervals, while a larger size is more accurate for data collected in longer intervals. Small rolling windows generate more volatile coefficients, while longer windows yield smoother estimates. In order to find the window size yielding the best predictions, a common approach is to analyse the Root Mean Square Error (RMSE). The purpose of forecasting is to minimize the forecast error, that is the difference between actual and predicted values. The RMSE can be defined as:

$$RMSE = \sqrt{\sum_{i=1}^{n} \frac{(\widehat{y_i} - y_i)^2}{n}} \quad (18)$$

where $\widehat{y_i}$ indicates the forecasted values and $y_i$indicated the actual values. Several regressions are run with different window sizes, ranging from two to twelve years. The model with the lowest RMSE is chosen as it possesses the highest forecast accuracy. Table 10 in Section 7.2.1 presents the chosen optimal window sizes for the countries in the sample.

As discussed in the hypothesis development, the first hypothesis tests whether oil prices have an effect on stock market returns. The null hypothesis is that the true value of the oil coefficient is equal to zero (i.e., no oil effect), whereas the alternative hypothesis states that the value of the oil coefficient is not equal to zero (i.e., oil effect). In order to test this hypothesis, the confidence interval approach is adopted with a 95% confidence level. The confidence interval represents the range of values within which the parameter under consideration is likely to reside (Brooks, 2008). In our case, we investigate whether the hypothesised value of the oil coefficient (i.e., zero) lies within the boundaries. In order to obtain the confidence interval, the estimates of the oil coefficient ($\hat{\beta}$) and its standard error (SE) are generated. Both lower and upper bounds can be calculated using the following formula, respectively:

$$\hat{\beta} - 1.96 * SE\left(\hat{\beta}\right) \;\; and \;\; \hat{\beta} + 1.96 * SE\left(\hat{\beta}\right) \tag{19}$$

where 1.96 is the critical value of the chosen 95% confidence interval. By plotting the bounds together with the coefficient estimates, it becomes clear how the oil coefficient has changed over time and at which points it is significant.

6.3.2 Underreaction to Oil Price Changes

In the second part of the analysis, we test whether investors underreact to oil prices. If stock markets are assumed to be efficient, one would expect to witness an instant reaction to oil price fluctuations, as new information is immediately incorporated in the stock market. However, previous empirical studies, conducted by amongst others Driesprong et al. (2008) and Fan and Jahan-Parvar (2011), find a delayed reaction of investors to oil price changes of several trading days, signalling market inefficiency.

The methodology used for our second hypothesis is slightly different compared to the methodology used in our first hypothesis. Here, we use a single static OLS regression over the last 11 years, as opposed to a multiple rolling-window regression considering the whole sample period. The reason for this 11-year period, between January 2010 and December 2020, as opposed to considering the whole period, is that the speed and way investors receive information has changed over time. Therefore, testing for underreaction twenty years ago is not representative for today’s fast-paced digital world. The choice of a static regression over a rolling regression is because we are no longer examining the relation between oil prices and stock returns over time, but how investors’ behaviour affects stock market performance and whether they underreact to oil price changes. Last, a single regression is conducted instead of a multiple regression. The industrial production index has been dropped as we are no longer attempting to control for changes in global growth, but we solely aim to research the underreaction behaviour.

To capture the underreaction, lags of several trading days are added between the monthly oil price observations and stock returns. Therefore, a new monthly oil price series with a delay of X number of trading days is constructed, using daily data of the WTI future prices. Since we do not know the duration of the delayed reaction in our sample, we try several lag lengths, namely lags of zero (i.e., contemporaneous relationship), two, four, six, eight and ten days. Regressions have been estimated for each country for each lag, according to the following equation:

$$r_t^{stock} = \alpha + \beta * r_{t-x}^{oil} \quad where\ x = 0; 2; 4; 6; 8; 10 \tag{20}$$

where r_t^{stock} represents the monthly return on the stock indices at time t and r_{t-x}^{oil} represents changes in WTI future prices, lagged *x* days. Assuming there is a delayed reaction, the inclusion of a lag is expected to increase the explanatory power of the regression as the model captures more of that delayed reaction. However, after a certain lag, the explanatory power is likely to decrease again, as by then the information has diffused amongst (the majority of) investors. To find the optimal lag, we compare the explanatory power of the regressions across lags, represented by the R^2. The R^2 is a statistic used to measure the goodness-of-fit of the regression model. This measure indicates the percentage in the dependent variable that the independent variable is able to explain and takes values between 0 and 1. Values close to 1 indicate that the model explains a high degree of variation in the dependent variable, whereas the opposite holds for values close to 0. Therefore, the higher the R^2, the better the regression fits the data (Brooks, 2008). Moreover, the t-values are taken into consideration to determine the significance of the results. The lag recording the highest $R^2$and most significant t-values indicates the day with the strongest oil effect and thus signals when the largest share of information regarding oil prices is incorporated in the stock market.

7. Empirical Results

7.1 Results of Stationarity Tests

Before testing our hypotheses, we determined whether the series are stationary. As discussed in Section 5.1, stock indices prices, oil futures prices and industrial production values were transformed by taking the natural logarithms of the returns at each point in time. Therefore, stationarity tests are conducted on both the original and transformed series. Starting with the WTI futures, both ADF and KPSS tests show that the original series is not stationary. The ADF test fails to reject the null hypothesis of non-stationarity with a p-value higher than the critical value of 0.05. Similarly, with t-statistics significantly higher than 0.463, the KPSS test rejects the null hypothesis of stationarity. These findings are confirmed when plotting the WTI futures prices, where seasonal fluctuations in the data can be identified. This is not surprising as oil prices are highly affected by changes in supply and demand levels and have sharply fluctuated over the sample considered (see Figure 2). Furthermore, non-stationarity can also be inferred from the ACF plot, where the autocorrelation values are slowly decreasing and remain well above zero for 100 lags (see Figure B1 in Appendix B). Converting the original series into logarithmic returns removes seasonality, thus making the series stationary, as confirmed by the results of both ADF and KPSS tests presented in Table 7.

Table 7: Results of the Stationarity Tests for the Oil Futures Variable

Part A: Original Series				
	ADF Test		KPSS Test	
	t-value	*p-value*	*t-value*	*p-value*
WTI	-2.3103	0.1687	1.5490	0.0100
Part B: Transformed Series				
	ADF Test		KPSS Test	
	t-value	*p-value*	*t-value*	*p-value*
WTI	-11.0796	0.0000	0.0876	0.1000

Similar to the oil futures, the industrial production index shows a trend, mostly upward, throughout the time period under consideration (see Figure 4). Both ADF and KPSS tests show that the original series is not stationary. The ADF test fails to reject the null hypothesis with a p-value of 0.4322. Similarly, the KPSS test rejects the null hypothesis of stationarity with a t-statistic significantly higher than 0.463. Non-stationarity can also be detected from the autocorrelation plot, where the values slowly decrease over 100 lags, but remain well above zero (see Figure B2 in Appendix B). Once the ADF and KPSS tests are run on the transformed series, both approaches indicate stationarity in the series of logarithmic returns, as reported in Table 8.

Table 8: Results of the Stationarity Tests for the Industrial Production Variable

Part A: Original Series				
	ADF Test		KPSS Test	
	t-value	*p-value*	*t-value*	*p-value*
IPI	-1.6982	0.4321	2.2881	0.0100
Part B: Transformed Series				
	ADF Test		KPSS Test	
	t-value	*p-value*	*t-value*	*p-value*
IPI	-14.8243	0.0000	0.3629	0.0931

Finally, Table 9 reports the findings of the tests conducted on the stock market series. It is not surprising that both approaches conclude on the non-stationarity of the original price series. Index prices are highly influenced by market forces and fluctuate as a consequence of various factors, such as economic growth, industrial production and GDP or inflation. Therefore, they cannot be stationary. This can also be concluded by looking at Figure 5, where the dynamics of the stock market indices are presented. They show both up- and down-ward trends, which confirms the non-stationarity of the data. The autocorrelation plots also signal non-stationarity in the countries' stock indices (see Figures B3 in Appendix B). Transforming the price series into logarithmic returns induce stationarity in the data, as can be seen from the results reported in Table 9.

Table 9: Results of the Stationarity Tests for the Stock Market Variables

	Part A: Original Series				Part B: Transformed Series			
	ADF Test		KPSS Test		ADF Test		KPSS Test	
	t-value	*p-value*	*t-value*	*p-value*	*t-value*	*p-value*	*t-value*	*p-value*
United States	1.6863	0.9981	2.1205	0.0100	-20.5102	0.0000	0.1038	0.1000
Mexico	-0.3323	0.9208	2.1473	0.0100	-20.1010	0.0000	0.3597	0.1000
Canada	-0.5403	0.8839	2.4109	0.0100	-19.1186	0.0000	0.0684	0.1000
Germany	-0.0450	0.9546	2.1912	0.0100	-20.1132	0.0000	0.0715	0.1000
China	-1.8133	0.3739	1.6200	0.0100	-5.6200	0.0000	0.2033	0.1000
Norway	-0.8967	0.7891	2.0486	0.0100	-16.9562	0.0000	0.0463	0.1000
Saudi Arabia	-1.8374	0.3621	1.0186	0.0100	-15.0318	0.0000	0.0938	0.1000
Brazil	0.9620	0.9938	1.9035	0.0100	-4.3911	0.0000	0.3752	0.1000
Russia	0.0816	0.9648	1.5693	0.0100	-14.5054	0.0000	0.1070	0.1000
Jordan	-2.0489	0.2656	1.2163	0.0100	-8.6399	0.0000	0.2589	0.1000
Hong Kong	-1.1964	0.6752	2.3560	0.0100	-21.0472	0.0000	0.2816	0.1000

7.2 Results of Empirical Analysis

7.2.1 Oil Effect

This section presents the empirical results of the rolling-window regressions estimated in order to test the first hypothesis of the research, that is, whether there is an oil effect. As discussed in Section 6.3.1, the performance of the rolling regression is sensitive to the choice of window size. In order to find the optimal window size, regressions with windows between two and twelve years are run, and the window with the lowest RMSE is chosen. The RMSE measures the prediction accuracy of the model and a low RMSE indicates that the model's predictions are accurate. Table 10 contains the RMSE of the ten countries for different window sizes, where the numbers in bold indicate the lowest RMSE for each country.

Table 10: RMSE per country over different window sizes

	RMSE					
Years	2	3	5	8	10	12
Window	24	36	60	96	120	144
United States	0.0436	0.0437	0.0421	**0.0420**	0.0428	0.0432
Canada	0.0513	0.0463	0.0402	**0.0399**	0.0420	0.0407
Norway	0.0641	0.0624	0.0590	0.0581	0.0611	**0.0576**
Russia	0.0818	0.0763	0.0667	0.0677	0.0725	**0.0523**
Brazil	0.1152	0.1121	0.0930	0.0782	0.0741	**0.0635**
China	0.1171	0.1023	0.0877	**0.0815**	0.0845	0.0842
Saudi Arabia	**0.0658**	0.0670	0.0676	0.0714	0.0708	0.0717
Mexico	0.0758	0.0710	0.0684	0.0624	0.0619	**0.0510**
Jordan	0.0481	**0.0470**	0.0480	0.0504	0.0512	0.0520
Hong Kong	0.0777	0.0754	0.0699	0.0696	0.0717	**0.0680**

Following Equations 15 and 16, several regressions have been obtained for the countries in the analysis, with the purpose of estimating the oil coefficient over time. Figure 5 displays the plots of the oil coefficient over time, which illustrates the evolution of the relationship between the oil price and the stock index for each country. A 95% confidence interval is plotted around the oil coefficient to show whether and where the oil effect is significantly different from zero. A positive oil coefficient implies that an increase in the oil price results in an increase in stock prices, whereas the opposite applies in case of a negative oil coefficient. By plotting the estimated oil coefficient throughout the entire period, it can be observed that it is time-varying for all the countries in the sample.

The regression results presented in Figure 5 show how the contemporaneous relationship evolves over time for each country. In the case of the United States, the oil coefficient is primarily negative up to 2006, but only significantly negative in the periods 1999-2001. In 2008, it becomes significantly positive up to 2011, where the relationship becomes insignificant. In 2020, the relationship turns significantly positive again. Next, Mexico shows an overall mildly positive relationship between oil prices and stock market and we find a significantly positive relationship after 2006. In the case of Hong Kong, the oil coefficient is negative up to 2003, but only significant between 1998 and 1999. The oil coefficient turns positive in 2003 and significantly positive in 2009 for the rest of the sample period. Following, the coefficient for Canada is (partly significantly) negative up to 1999. Afterwards, the coefficient switches directionality in 1999, and becomes significantly positive in mid-2002. As far as China is concerned, the relationship between oil prices and the stock index is overall mildly positive, except before the year 2003, and only significant between 2005 and 2015. Next, looking at the Norwegian stock market, the oil effect is always positive during the whole period, and significantly positive after 2008. Norway has a relatively high oil coefficient between 2015 and 2016, where it almost reaches 0.4. In the case of Jordan, the oil coefficient is relatively volatile throughout the period and therefore only between 1996-1998 a significant (negative) relationship is detected. During the rest of the period, the relationship is insignificant, with some very short-term (positive) exceptions in mid-2001, 2010 and 2011. Similarly, the oil coefficient of Saudi Arabia is very volatile and therefore displays little and only brief periods of significance. A significantly negative relationship is found during the periods 2005-2007, while a significantly positive relationship is found between 2009-2011, 2015-2016 and after 2020. Regarding Brazil, the relationship is initially negative between 2003 and late 2008. Afterwards, the oil coefficient becomes positive, but significant only after 2010. Last, the oil coefficient of Russia is positively significant over the whole sample period and rather high and stable.

To conclude, all countries show an oil effect at least some point in time, meaning the null hypothesis of no oil effect can be rejected. Some common characteristics of the relationship can be observed from the results. First, in line with the expectations, the oil-exporting countries, being Mexico, Brazil, Russia, Norway and Canada, are almost entirely significantly positive. We observe that the positive relationship for two least developed (and exporting) countries, Mexico and Brazil, is less pronounced compared to Russia, Norway and Canada. As the demand for oil is growing at a higher rate for an emerging country as opposed to a more developed country, the positive effect of an increase in oil income as an exporter is partly offset by the increase in oil cost for production. Only for oil-exporter Saudi Arabia no clear relationship can be established. Saudi Arabia, together with oil-importer Jordan, show the least significance, which can be attributed to their strong dependence on oil, combined with high political instability. The oil-importing countries show less homogeneous results, possibly caused by exogenous factors.

Figure 5: Rolling Beta Estimates per Country

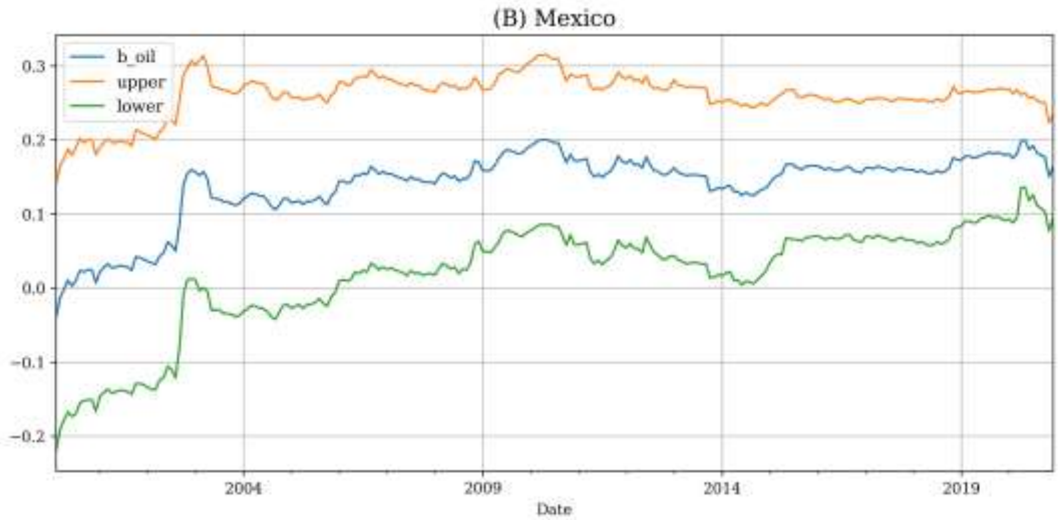

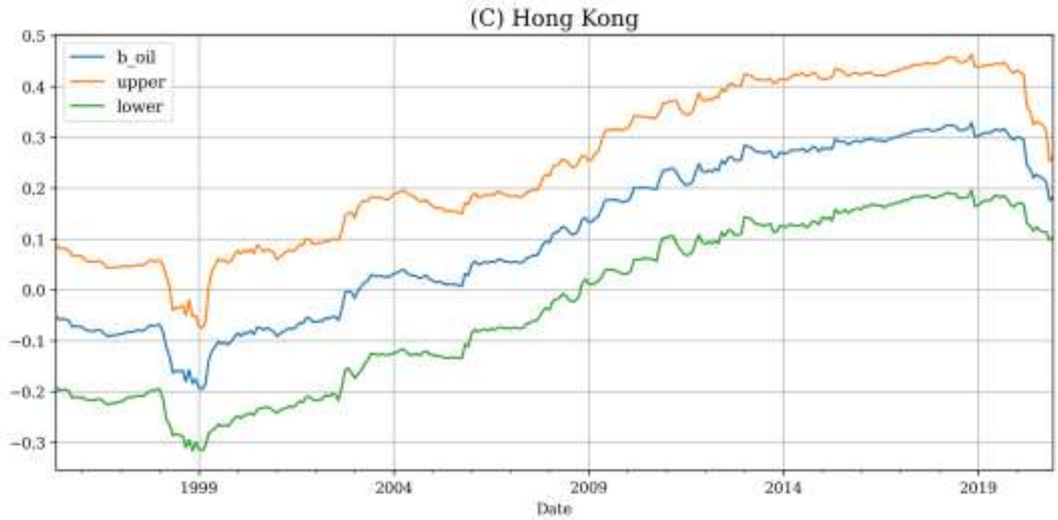

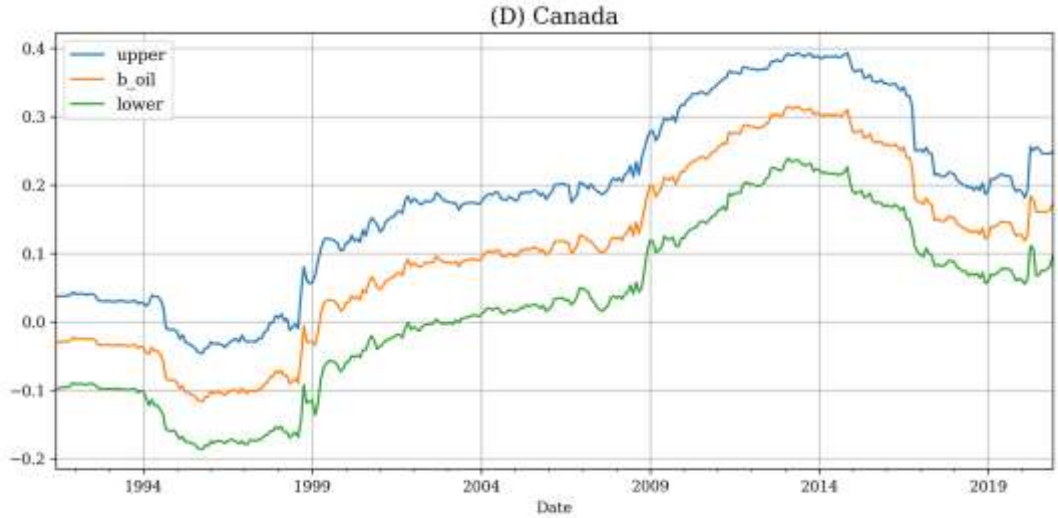
(D) Canada
upper
b_oil
lower
0.4
0.3
0.2
0.1
0.0
−0.1
−0.2
1994
1999
2004
2009
2014
2019
Date

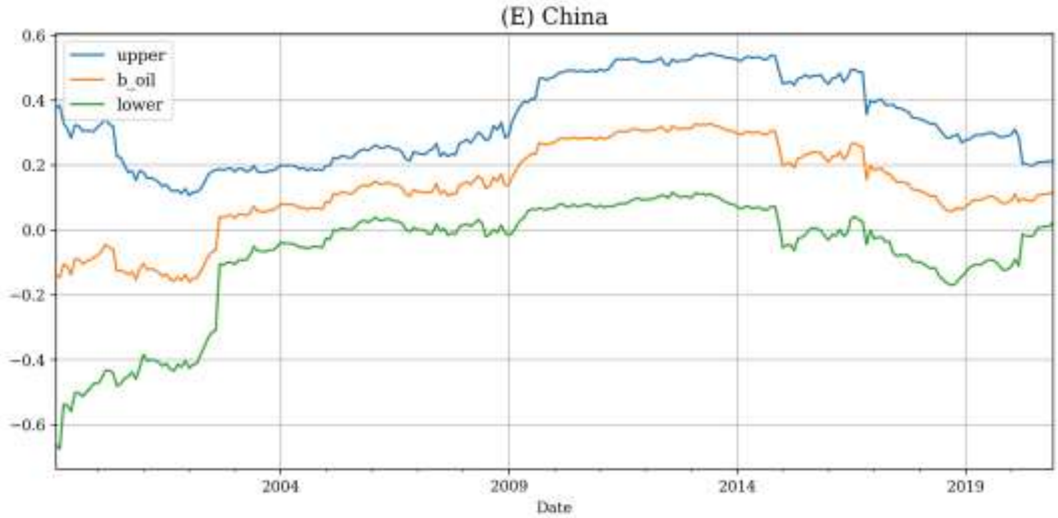
(E) China
upper
b_oil
lower
0.6
0.4
0.2
0.0
−0.2
−0.4
−0.6
2004
2009
2014
2019
Date

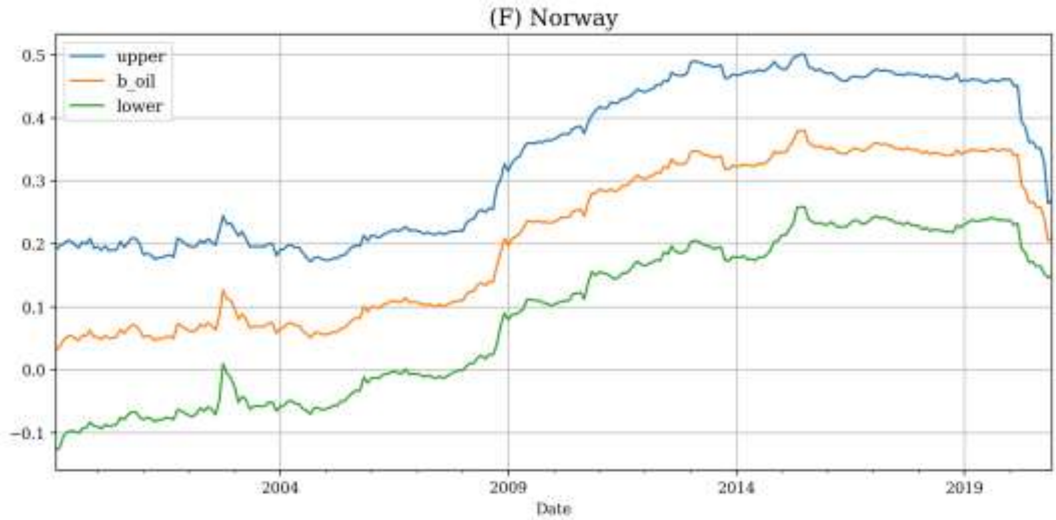
(F) Norway
upper
b_oil
lower
0.5
0.4
0.3
0.2
0.1
0.0
−0.1
2004
2009
2014
2019
Date

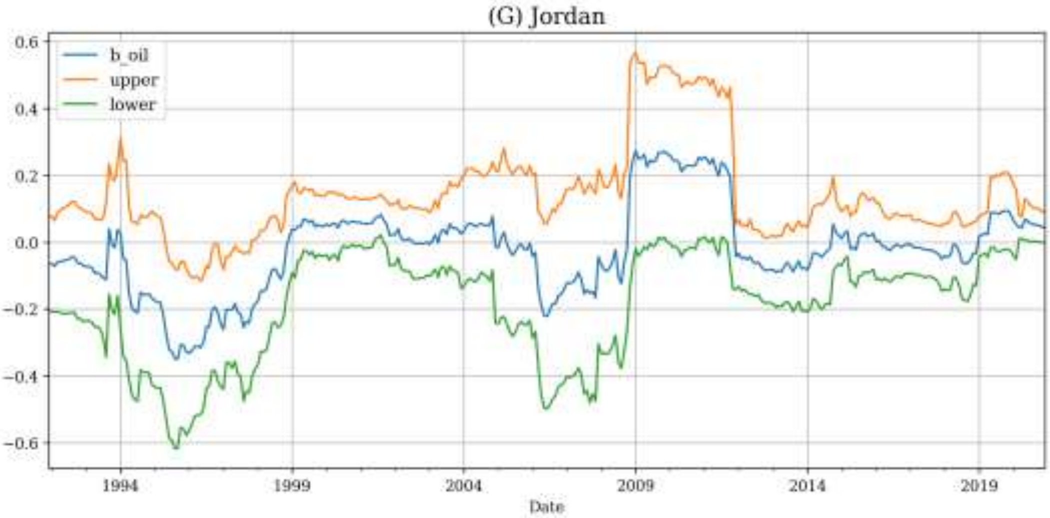
(G) Jordan
b_oil
upper
lower
0.6
0.4
0.2
0.0
−0.2
−0.4
−0.6
1994
1999
2004
2009
2014
2019
Date

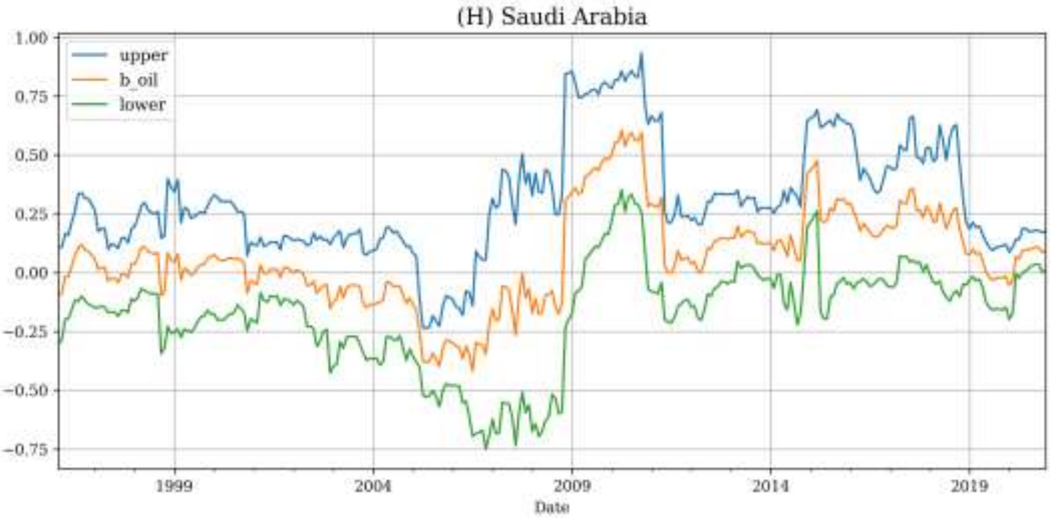
(H) Saudi Arabia
upper
b_oil
lower
1.00
0.75
0.50
0.25
0.00
−0.25
−0.50
−0.75
1999
2004
2009
2014
2019
Date

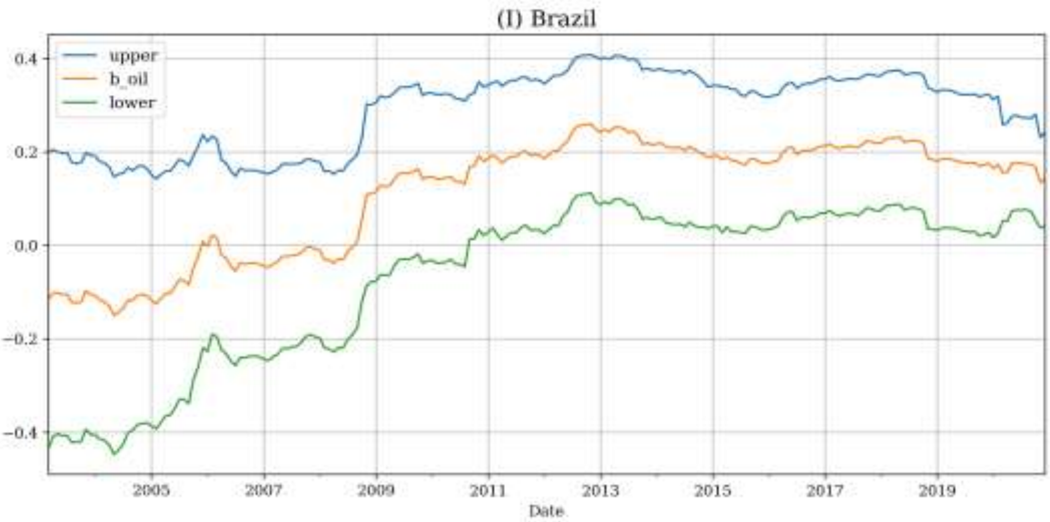
(I) Brazil
upper
b_oil
lower
0.4
0.2
0.0
−0.2
−0.4
2005
2007
2009
2011
2013
2015
2017
2019
Date

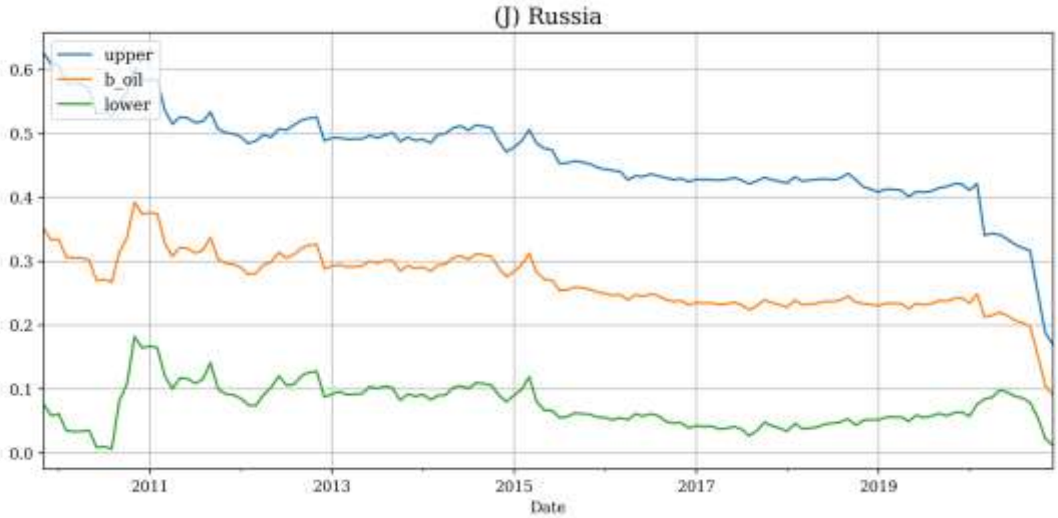

7.2.2 Underreaction to Oil Prices

Next, we test whether investors react with a delay to oil price changes. To do so, we examine whether the oil effect increases after including lags of several trading days. An increase of oil effect using the lagged oil price series would signal investors underreaction and can be explained by the gradual information diffusion hypothesis introduced by Hong and Stein (1999). In order to test for underreaction, we compare the explanatory power of the regressions across lags. Figure 6 plots the explanatory power, R^2, as a function of the number of trading days for each country in the sample. Following Equation 20, regressions are run with lags of zero, two, four, six, eight and ten days. The optimal lag size varies across the countries in the sample. In the case of the United States, Hong Kong and Brazil, the explanatory power is the lowest after two trading days, peaks at the lag of eight trading days, and then decreases as the lag size further increases. Similarly, the explanatory power is the highest at lag eight for Canada, Norway, Mexico and Saudi Arabia. Conversely, in Jordan and Russia, the oil effect is the strongest at lag zero, thus signalling little delayed reaction. Finally, in the case of China, the regression with two lagged days reports the lowest explanatory power, which then peaks at a lag of four trading days.

Figure 6: R-squared as a Function of Lags

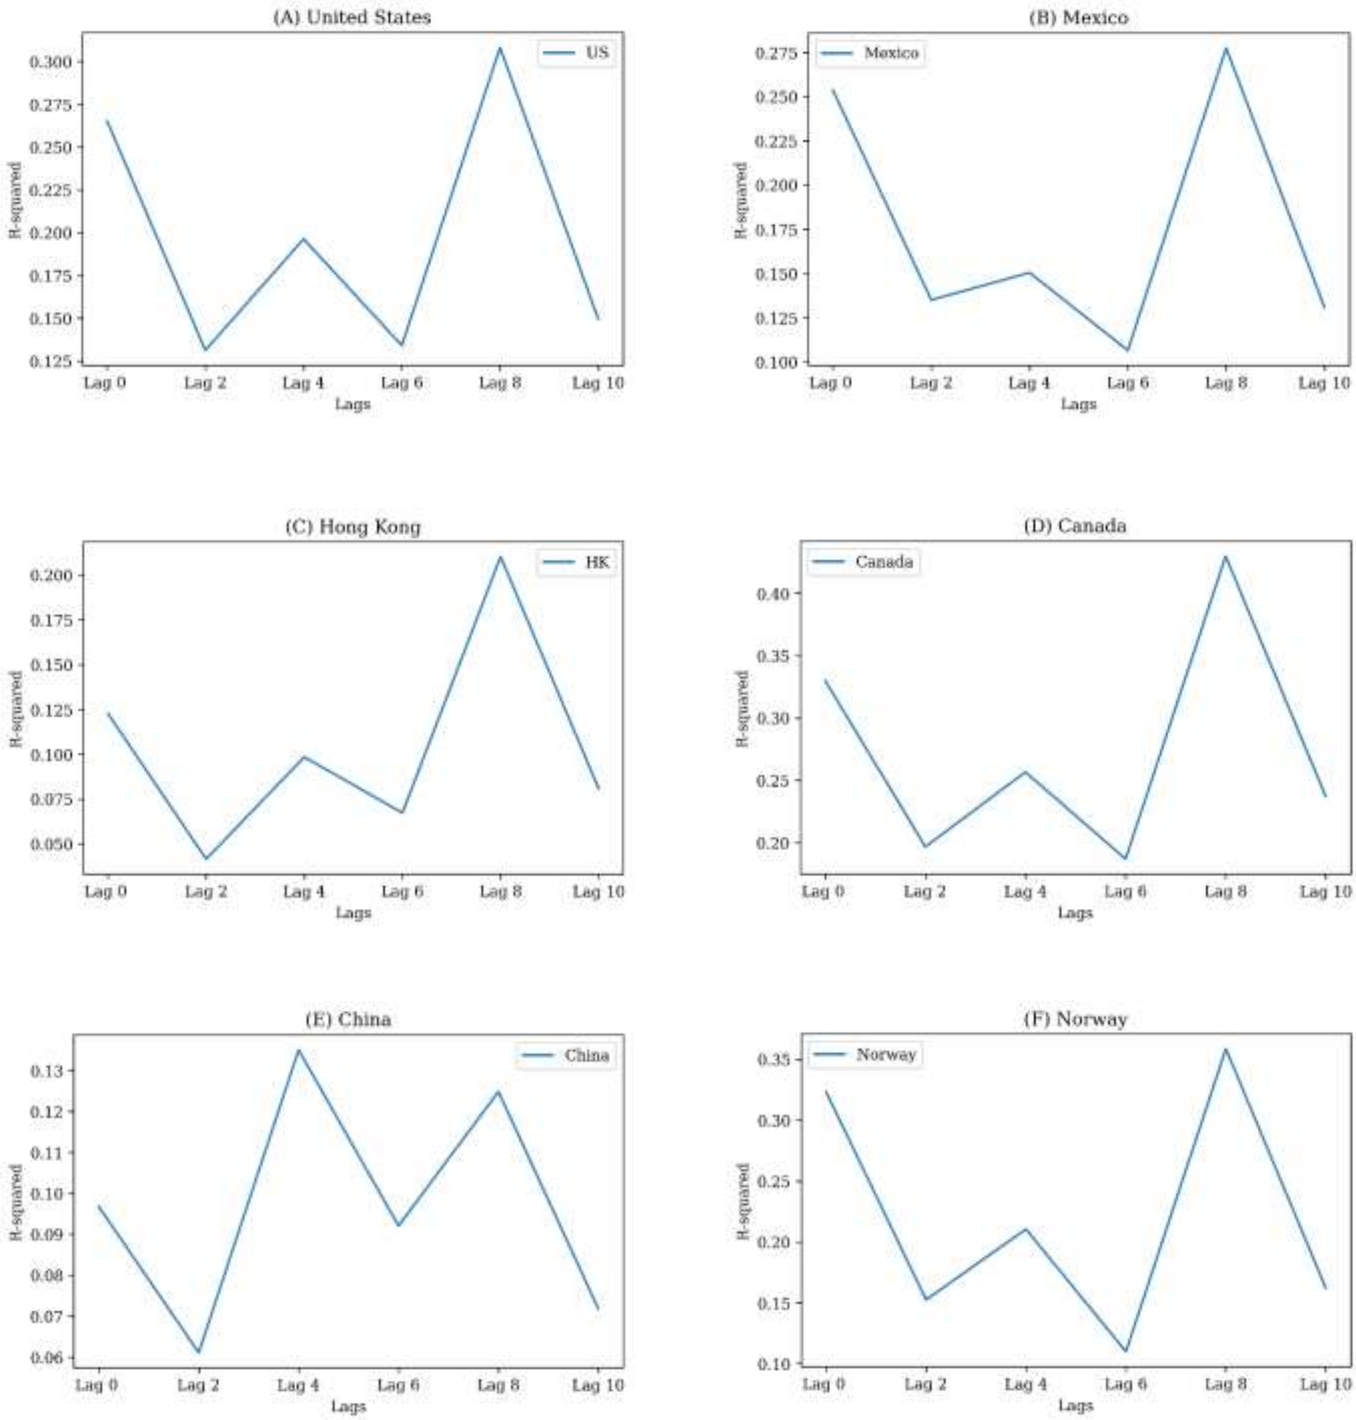

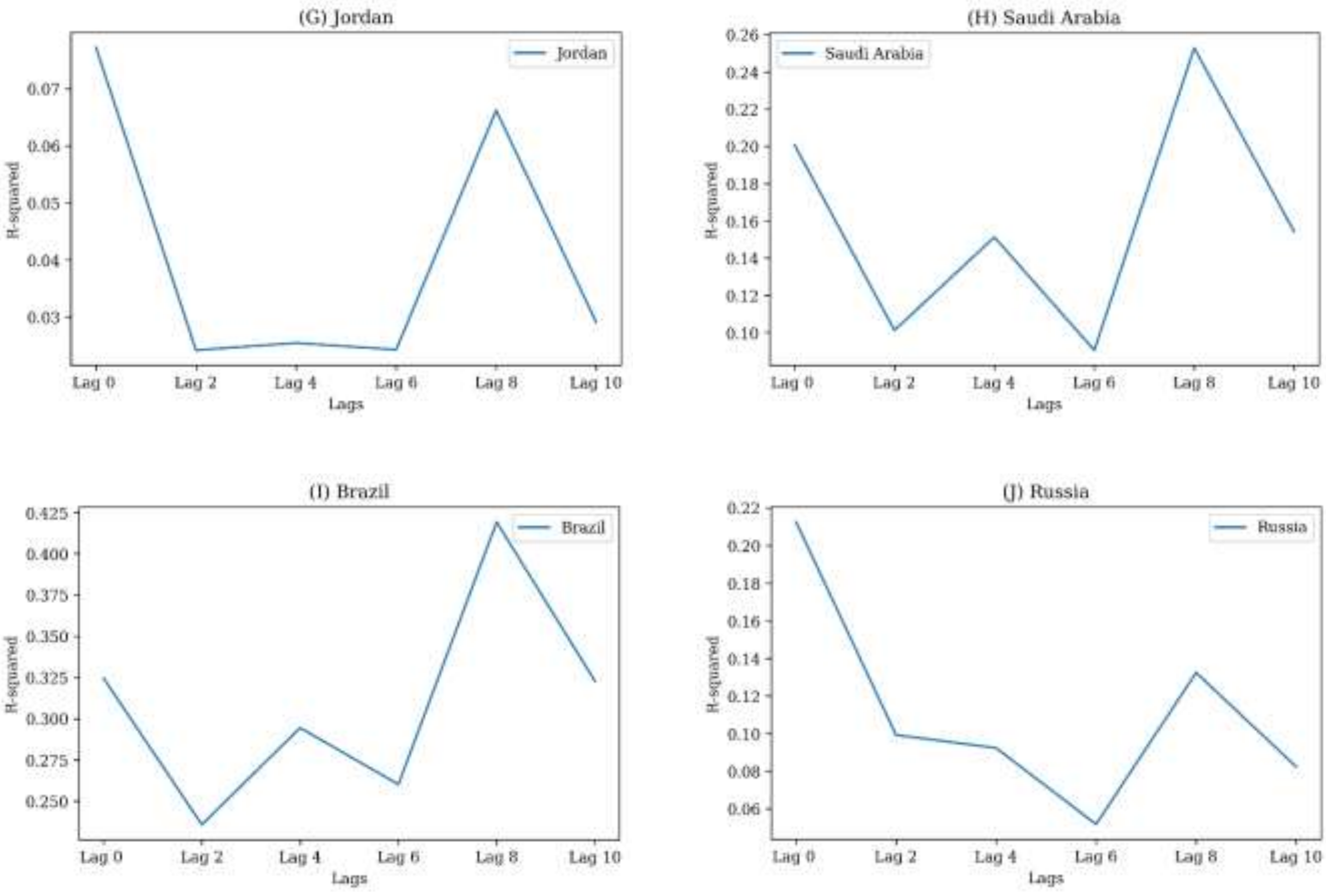

Table 11 contains the estimation results, namely the t-values and R^2, based on Equation 20. As discussed above, lag eight shows the highest explanatory power overall, followed by lag zero, while lag two and six report an overall relatively low explanatory power. These outcomes are further strengthened when comparing the significance level across lags. Starting with the United States, Hong Kong and Brazil, we find a strong relationship between monthly oil returns and stock market returns both at lag zero and lag eight. However, the results of the eight-day lag period indicate the strongest oil effect as the R^2 is larger and t-values are more significant (at a 95% significance level). Next, in the case of Canada, Norway, Mexico and Saudi Arabia, the size of the t-values and R^2 is largest at lag eight and lowest at lag six. For Jordan and Russia, we find the strongest relationship between oil returns and stock market returns at lag zero. The statistics are the lowest for lag two for both countries, with the t-value becoming insignificant in Jordan. Finally, the t-statistic for China indicates that lag four has the highest explanatory power.

Overall, for the majority of the countries, the explanatory power spikes at day eight, signalling that it takes around eight trading days before information about oil price shocks is fully incorporated in the market. These findings show evidence that investors in several countries underreact to news regarding oil price changes with a delay of several trading days. This delayed reaction can lead to stock return predictability and therefore can have implications for investors. To this end, trading strategies based on both the oil effect and the detected underreaction are developed and discussed in details in Section 10.

Table 11: Regression results with different lags between monthly stock returns and oil price changes

	Lag 0		Lag 2		Lag 4	
	t-value	R-squared	t-value	R-squared	t-value	R-squared
United States	4.4418	0.2652	3.2286	0.1312	4.1051	0.1963
Mexico	4.8419	0.2536	3.2811	0.1349	3.4941	0.1503
Hong Kong	3.1058	0.1226	1.7289	0.0415	2.7444	0.0984
Canada	6.0524	0.3298	4.1105	0.1967	4.8809	0.2567
China	2.7180	0.0967	2.1194	0.0611	3.2814	0.1350
Norway	5.7248	0.3233	3.5238	0.1525	4.2892	0.2105
Jordan	2.4040	0.0772	1.3105	0.0242	1.3429	0.0255
Saudi Arabia	4.1647	0.2008	2.7881	0.1012	3.5059	0.1512
Brazil	5.7579	0.3245	4.6142	0.2358	5.3668	0.2945
Russia	4.3152	0.2125	2.7576	0.0992	2.6492	0.0923

	Lag 6		Lag 8		Lag 10	
	t-value	R-squared	t-value	R-squared	t-value	R-squared
United States	3.2679	0.1340	5.5443	0.3082	3.4824	0.1495
Mexico	2.8654	0.1063	5.1486	0.2775	3.2209	0.1307
Hong Kong	2.2305	0.0673	4.2842	0.2101	2.4609	0.0807
Canada	3.9826	0.1869	7.2111	0.4297	4.6305	0.2371
China	2.6448	0.0920	3.1369	0.1248	2.3106	0.0718
Norway	2.9182	0.1099	6.2115	0.3586	3.6551	0.1622
Jordan	1.3105	0.0243	2.2126	0.0662	1.4377	0.0291
Saudi Arabia	2.6199	0.0905	4.8324	0.2528	3.5464	0.1542
Brazil	4.9268	0.2602	7.0584	0.4192	5.7354	0.3228
Russia	1.9400	0.0517	3.2452	0.1324	2.4883	0.0823

8. Discussion

8.1 Oil effect

The results in Section 7.2.1 present how the contemporaneous relationship between oil prices and stock markets per country evolves over time. After evaluating the results, we can clearly observe that the relationship differs per country and evolves over time. In this section, the results are discussed and compared to the expectations discussed in the literature and theoretical review, while considering exogenous effects potentially influencing the relationship.

The time-varying relationship between the U.S. stock market index and WTI prices starts overall negative, which is in line with the expectation that oil-importing countries show a negative relationship. This initial negative relationship turned positive during the financial crisis of 2008/2009. This can be explained by the fact that the crisis caused stock markets to enter bearish territories, while simultaneously causing the oil prices to largely decrease. Thus, both oil prices and stock prices dropped together, leading to a positive relationship. Since the financial crisis, which made global financial markets more interdependent, the time-varying relationship between WTI prices and the U.S. stock market index has strengthened and remained positive. However, the relationship began to decline in 2011, a period associated with continuing unrest in the Middle East during 2011-2014 and the Syrian civil war. These events triggered positive oil-market specific price shocks, and stock markets responded negatively to such news, causing the correlation to decline. The downturn in the relationship continued until 2018, as a consequence of, among others, the slowdown in the Chinese economy, Iran’s return to international trade, the decline in global demand, and the rivalry between the United States and Saudi Arabia for the control of the oil market (Youssef & Mokni, 2019). In September 2019, the United States became a net petroleum exporter, which explains the increase in positive relationship (U.S. Energy Information Administration, 2020b). In spring 2020, oil prices collapsed amid the Covid-19 pandemic and economic slowdown. In parallel, the U.S. stock market plummeted by 20% in three days (Imbert, 2020) . This event causes the coefficient to increase again and the relationship to become significantly positive.

China is the world’s largest energy consumer but also ranks within the top five of oil producers (U.S. Energy Information Administration, 2020a). Yet, China is still a net oil-importer due to its enormous demand for oil to sustain its fast-growing economy. Figure 5.E shows that the oil coefficient fluctuates around zero in the beginning. This implies that the oil price is not that correlated with the stock market, in particular in the period leading up to 2009 as it is mainly insignificant. This is in line with Nguyen and Bhatti (2012), who claim that the Chinese stock market index is not affected by oil prices during the period 2000–2009. A plausible explanation is the fact that China had a strong and stable economic growth during that period and thus the country was able to absorb the changes in oil prices. Furthermore, China's stock market was mainly driven by internal factors until the global financial crisis. After the crisis, the urbanisation and industrialisation of China led its stock market to become more influenced by international factors, which is shown by an increase in the coefficient, which became significantly positive in 2009. Overall, the rather low correlation implies that the Chinese stock market is relatively 'secure' against oil market-based risk. This is in line with Broadstock and Filis (2014), that argue that the Chinese stock market is rather resilient to oil price shocks.

Hong Kong is an oil-importing region. Therefore, the obtained positive relationship contradicts the expectation. This positive oil effect on the Hong Kong stock market can be attributed to the fact that a large number of Chinese companies are listed on the Hong Kong stock exchange and investors expect the Chinese economy to perform well, even in case of higher oil prices. Therefore, investors still expect capital to flow into Hong Kong, also in case of higher oil prices. Thus, the positive expectation effect of China's fast-growing economy may be greater than the negative effect of the precautionary demand-driven effect. These results are in line with the findings of, among others, Lin et al. (2010).

The linkage between oil and stock markets in oil-exporting countries are positive for most of the sample period. This can be explained by the fact that global economic activity and demand for oil generate upward pressures on the price of crude oil, which in turn leads to higher stock prices for oil-exporting countries. Thus, oil and stock prices move in the same direction. This positive relationship can most strongly be observed for Canada, Norway and Russia. In the case of Canada, the relationship between the stock market and WTI futures is positive only after 1999. Before 1999, Canada's status as leading oil-exporter was not as pronounced, which explains the negative coefficient. However, Canada is currently the world's fourth-largest exporter of oil, thus explaining the positive and relatively high oil coefficient (Government of Canada, 2020). Similar positive relationships hold for Russia and Norway. Russia is the second largest oil-exporting country in 2019 and its stock exchange is dominated by oil & gas companies, representing approximately 60% of the total market capitalization (U.S. Energy

Information Administration, 2017). The Russian stock market registers high volatility (as shown in Figure 5.J), indicating that the market is exposed to oil price risk, which is in line with its strong dependency on oil. Norway is Europe's largest producer of oil and gas. However, the oil sector is only 20 % of the weight of its stock index, meaning that the remaining 80% is less or possibly inversely correlated to the oil price (Norsk Petroleum, 2020). This makes Norway less sensitive to oil price changes compared to Russia, which can be observed in the lower volatility of the Norwegian stock index and the overall lower estimated oil coefficient.

Mexico and Brazil are also categorized as oil-exporting countries. They show a positive relationship, which is however not as strong and significant as in the case of Canada, Russia and Norway. This can be attributed to a difference in the level of development. As the demand for oil is rapidly growing for an emerging country, the positive effect of an increase in oil income as an exporter is partly offset by the increase in oil cost for production (Alekhina & Yoshino, 2018). Therefore, the relationship is not as positive as for developed countries such as Norway, as the demand for oil in developed countries is often holding steady or declining slightly.

Saudi Arabia is an OPEC country and is one of the largest oil producers in the world, as it controls 17% of the global oil reserves (OPEC, 2019). This suggests that its economy is heavily dependent on oil. Due to its strong dependency, the stock market can be very sensitive to oil price changes. Thus, the coefficient is very volatile and often insignificant. A significant negative relationship is found in 2005/2006, which contradicts the expectation of an oil-exporting country. A reason for this negative relationship can be that the oil prices plummeted in 2005, while at the same time Saudi Arabia became a member of the World Trade Organization, causing an increase in foreign investment which boosted the economy. A positive significant relationship is found after the financial crisis in 2008/2009. The period hereafter is characterised by increased political instability during the aftermath of the Arab uprising in late 2010, causing extreme variability in both oil and financial markets and resulting in an insignificant coefficient. The relationship briefly turned significantly positive after the oil price plunge of 2014/2015, which considerably negatively impacted equity returns in Saudi Arabia.

Similarly to Saudi Arabia, Jordan is heavily dependent on oil and largely exposed to the developments in the crude oil market. However, it is a net-oil importer and thus a negative relationship is expected. The relationship shows large volatility and is overall insignificant, except during 1996-1998 where a significantly negative relation is detected. Thereafter, the results remain insignificant, with some short-lived positive exceptions in mid-2001, 2010 and 2011. These positive, yet insignificant, results are in line with Berument et al. (2010) and Al-Fayoumi (2009). They suggest that the Jordanian stock market is sensitive to oil price variations, as it shows asymmetric responses to variations in crude oil prices, both positive and negative.

Overall, our findings confirm our first hypothesis that there is a contemporaneous relationship between oil price and stock returns. It becomes clear that the relationship does not remain constant but behaves heterogeneously over different periods in time, which is in line with other research (i.e., Bahr & Nikolova, 2010; Antonakakis et al., 2013). For the oil-exporting countries, besides Saudi Arabia, we find a positive and almost entirely significant relationship, which is in line with expectations. Noteworthy is that the positive relationship for the two least developed (exporting) countries, Mexico and Brazil, is less pronounced compared to Russia, Norway and Canada. As the demand for oil is generally high in emerging countries, the positive effect of an increase in oil income is partly neutralized by the increase in oil cost for production. The results of both Saudi Arabia and Jordan show little significance and the coefficient is highly volatile, which can be attributed to their high dependence on oil and political instability. For the oil-importing countries, the relationship is often more complicated and not completely significant. We attribute these heterogeneous findings to the notion that the oil market is subject to exogenous forces and to the asymmetric relationship between oil prices and stock markets.

8.2 Underreaction to Oil Price Changes

As reported in the results in Section 7.2.2, we find evidence of a delayed reaction of investors of different trading days depending on the country. In the case of the United States, Hong Kong, Canada, Norway, Mexico, Saudi Arabia and Brazil, investors react most strongly with a delay of eight trading days to oil prices. For China, the delayed reaction is the strongest after four days. In the case of Jordan and Russia, there seems to be the least underreaction, as the strongest oil effect is found in lag zero. This means that for Jordan and Russia the oil effect is rather short-lived and new information is quickly incorporated in the stock prices. Our results are in line with previous empirical research, amongst others Driesprong et al. (2008) and Fan and Jahan-Parvar (2011), who both find a delayed reaction of investors.

This underreaction effect is often attributed to the bounded rationality of investors. Their ability to make rational decisions is challenged by their emotions and faulty behavioural biases (Hong & Stein, 1999). The delayed reaction challenges the notion of market efficiency. If markets are efficient, one would assume that news is immediately incorporated in the stock markets. Theoretically, this should apply to oil prices, as these are publicly available information and can be observed in real time and at no cost. However, our results indicate that it takes time before such information is available to all investors and is fully reflected in the stock market. This anomaly can be explained by the gradual diffusion hypothesis formulated by Hong and Stein (1999). Information gradually diffuses across the investors' community as traders underestimate the direct effect of oil shocks on the stock markets. This delayed reaction of investors can be attributed to the fact that they find it difficult to correctly assess the impact of oil prices on the value of stocks and on the general economy. The difficulty might lie in the fact that oil prices influence stock returns through various channels with sometimes conflicting directions. For instance, oil prices can affect a country's future cash flows either positively or negatively, depending on whether the country is an oil-consumer or oil-producer. In addition, traders generally focus solely on the market they specialize in and fail to capture potential valuable information coming from other markets in a timely manner. This can be due to the fact that paying attention to news coming from different markets is subject to time-budgeting constraint (Fan & Jahan-Parvar, 2011).

This underreaction behaviour has economic consequences, especially regarding the efficiency of financial markets. In an efficient market, stock prices adjust to new information immediately, suggesting that it is impossible to consistently obtain abnormal returns. On the contrary, when investors underreact to news, stock prices drift away from their intrinsic values and move less than the information justifies. Therefore, traders could exploit this anomaly by creating a trading strategy that earns on this market inefficiency.

9. Limitations and Future Research

It is plausible that a number of limitations have influenced the empirical results reported in this study. Limitations affecting the results of our first hypothesis, where we test for a contemporaneous relationship, mainly concern the methodology, namely the usage of a rolling-window regression. This methodology is chosen as it is one of the most common methods used in time series analysis to assess the stability of the parameters over time and it provides straightforward applications. The main drawback of a rolling window regression is that it is unable to capture immediate changes in the exposures as the estimated parameters depend on the length of the estimation window. Thus, the estimated oil coefficient cannot adapt quickly enough to changing oil price characteristics. Hence, the rolling-window approach works best under the assumption that the future is similar enough to the past, for example in periods with little volatility and in stable countries (quote). Since the period under analysis exhibits quite some volatility, especially after 2008, the established relationships should be interpreted with caution. For further research, we would recommend using the Kalman filter or a time-varying parameter vector autoregression. These models work better at handling sudden changes in oil prices and volatility in the stock market. Another limitation related to the dependence upon the length of the estimation window, is that the rolling-window regression is highly sensitive to the choice of the window size (Inoue & Rossi , 2012). Meaning, different window lengths lead to different forecast performances. To determine the optimal window length, we chose the size that minimizes the RMSE. There are, however, more sophisticated window optimization techniques that could be considered using, such as cross validation and weighted average of forecasts methods suggested by Pesaran and Timmermann (2007), the weighted least squares method of Anatolyev and Kitov (2007), or the AIC bandwidth selection rule of (Cai, 2007). Another possible source of error is a likely endogeneity problem. Namely, we cannot rule out that stock index returns have an impact on crude oil futures prices. When endogeneity between the dependent and independent variable is a natural result of their theoretical relationship, it becomes increasingly challenging to draw conclusions about the direction of causality and the regression results become difficult to interpret. Possible solutions to the endogeneity problem include adding an instrumental variable that is exogenous and uncorrelated with the dependent variable of the OLS model or using lagged dependent variables (Wooldridge, 2008). Despite the limitations of our method and possible biased results due to endogeneity, our findings do provide a good indication of the time-varying relationship and its directionality.

The main limitation regarding our second hypothesis, where we examine the underreaction behaviour of investors, is the fact that there lies some uncertainty in the reasoning. However, it is noteworthy that underreaction, and behavioural finance concepts in general, are difficult to define and measure empirically. The challenging aspect to studying behavioural variables in real world data is that they are difficult to isolate, and the precise psychological underpinnings of behaviour often remain an open question. For further research, an option to examine underreaction behaviour in the oil-stock market relationship would be to use an experimental model, as the data availability requires the existence of a restrictive environment (Almeida & Pereira Câmara Leal, 2015). In our case, to be able to quantify underreaction, we construct a simplified method. With this method we do not take into account the underlying psychological reasoning of investors. Despite the fact that we leave out the underlying psychological reasoning of investors, which is out of scope to our studies, the method does give an indication of underreaction to some extent as it shows that the explanatory power, and thus the oil effect, increases after several lagged days.

For further research, it would be interesting to examine the effects of oil price changes on stock returns across industries within countries. Also, going more in-depth into the asymmetric relationship between oil prices and stock market would be recommended to get a better understanding of the relationship. However, as these asymmetric effects are a research field of its own, we chose not to include this in our analysis.

10. Can we use oil prices to obtain abnormal returns?

Testing our two hypotheses has provided us with interesting insights. In the case of the first hypothesis, we find that oil prices do exert an effect on stock market returns. This directionality of the effect is not homogeneous and depends on many factors, among others on whether the country is oil-importing or exporting and on the underlying reasons causing the oil price change. In the case of the second hypothesis, we find that investors react with a delay to oil price changes. For the majority of the countries this delayed reaction takes up to eight trading days. Taking these findings into consideration, we are interested in determining whether they convey any exploitable financial conclusions. We therefore develop two trading strategies based on oil price changes. The purpose is to obtain some indication regarding the performance of an oil strategy, taking investors' underreaction into consideration.

We backtest the trading strategies over the period from January 2010 to December 2020, totalling a number of 132 observations for each country. This 11-year period is consistent with the period used to research underreaction, as we base our choice of lag on the delayed reaction of investors tested in the second hypothesis. More precisely, the lag with the highest explanatory power, in terms of R^2, for each country is used in the trading strategies. This is motivated by the fact that, as the strategies are based on the predictions of stock returns, the lag having the highest explanatory power should, ceteris paribus, yield the most accurate predictions. The trading strategies are tested on the stock markets of Mexico, Hong Kong, Canada, Norway and Brazil. This set of countries is chosen as, when testing the first hypothesis, we find a significant relationship between oil prices and stock returns over, at least, the last eleven years. Despite also finding a significant relationship for Russia over the 11-year period, Russia is excluded from the sample because the lag with the highest explanatory power is zero. Predictions on lag zero cannot be implemented in our trading strategies, as in our assumed real-world setting it is unrealistic to trade upon information at lag zero. The five included countries all show the highest level of delayed reaction at eight trading days, which can be traded upon. As predicted stock returns are later compared to the short-term U.S. Treasury bills, each stock market index is converted to USD.[4]

To construct the strategies, the expected stock returns are estimated for each country using a rolling-window regression, where the oil price is lagged on day eight. Unlike Equation

[4] The exchange rates for the different currencies have been retrieved from "*Investing.com*" for the period January 2010-December 2020.

16, we do not include the industrial production factor, as we aim to assess if abnormal returns can be obtained solely based on an oil effect. Thereby, as opposed to our previous analysis, we now consider simple returns instead of logarithmic returns. This is because linear returns are a better approach in the context of investments analysis and portfolio optimization. The optimal windows for the rolling regressions are chosen again based on the lowest RMSE, as presented in Table 12. where the numbers in bold indicate the chosen window size. Each regression is estimated using an initial number of observations equal to the window. Afterwards, the estimated parameters, being the intercept (α) and the oil coefficient (β), and the observed oil price changes in the previous eight days are used to create predictions of the stock market returns. The regression is then re-estimated every month using a sliding window to obtain forecasts of stock market returns. Based on these predictions, two trading strategies are developed.

Table 12: RMSE per country over different window sizes

Years	2	4	5	6	7
Window	24	48	60	72	84
Mexico	0.0636	**0.0635**	0.0657	0.0698	0.0740
Hong Kong	0.0493	0.0479	0.0485	**0.0449**	0.0453
Canada	**0.0404**	0.0426	0.0428	0.0435	0.0438
Norway	0.0545	**0.0520**	0.0534	0.0560	0.0584
Brazil	0.0942	0.0960	0.0974	0.1015	**0.0933**

The first strategy is a long-only strategy, where investors have two options, either invest in the stock market index or in a risk-free asset. As a proxy for this risk-free asset, the three-month U.S. Treasury Bill rate is used.[5] First, we verify whether the expected stock market return in the respective month exceeds the risk-free rate. If the expected return is higher than the risk-free rate, we fully invest in the respective stock index of the country. If the expected stock return is lower than the risk-free rate, we invest in three-month Treasury Bills. The second strategy is a long-short strategy, purely based on the sign of the expected stock returns. In case of positive expected returns, we long the respective stock index of the country, while in case of negative

[5] The data is retrieved from the database of the Federal Reserve Bank of St. Louis.

expected returns we short the index. The position for both strategies is rebalanced on a monthly basis, and trading costs are not considered for simplicity.

The risk and return features of these strategies are compared with a buy-and-hold portfolio, where we simply invest in the stock index of the country and hold the position during the entire period. To compare the strategies, the expected average returns, volatilities and the Sharpe Ratio (SR) are considered. The SR is defined as the ratio between the excess return and the standard deviation of each strategy, as indicated in the following formula:

$$SR = \frac{R_p - R_f}{\sigma_p} \tag{21}$$

The SR is a measure of risk-adjusted return and describes how much excess return the investors receive for taking extra risk (Munk, 2019). Generally, a SR greater than 1 is considered as acceptable by investors, while a ratio under 1 is defined as sub-optimal. Table 13 reports the annualized mean, standard deviation and the SR of the three trading strategies for each country. Figure 6 illustrates the expected returns for the two oil strategies and the buy-and-hold strategy for each country.

First, in the case of Mexico, Strategy 2 delivers a better performance than Strategy 1 and the buy-and-hold strategy, with a SR of 0.7577 and annualized expected return of 18.07%. As Figure 7.A shows, Strategy 2 generates higher returns consistently throughout the period, with a sharp increase in returns in 2020 due to the short position on the Mexican stock market index. Strategy 1 represents a less risky investment, as it partly invests in the risk-free asset. Interestingly, in the two years, we solely invest in the risk-free asset, as the predicted stock returns are consistently lower than the risk-free rate. Altogether, Strategy 2 shows the highest expected return and SR, thus possibly representing a profitable solution for investors interested in trading in the Mexican stock market.

Second, as far as Hong Kong is concerned, both Strategy 1 and 2 perform slightly better than the buy-and-hold strategy. Strategy 1 and Strategy 2 produce annualized expected returns of 11.65% and 13.52%, respectively. Figure 7.B shows that the buy-and-hold strategy outperforms the oil strategies between mid-2016 and beginning of 2020. The performance of Strategy 2 is relatively low before 2020 and this can be attributed to the fact that the model incorrectly predicts stock returns. However, in 2020, the return of Strategy 2 spikes due to shorting during a downturn of the index, as the expected returns are primarily negative.

Similarly, in 2020 the return of Strategy 1 increases and outperforms the buy-and-hold strategy through investments in the risk-free rate. Overall, Strategy 1 presents the most optimal investment opportunity in terms of risk-reward trade-off.

Third, in the case of Canada, both oil strategies deliver a considerably better performance than the buy-and-hold strategy. Holding the Canadian market index offers a rather low average return of 2.46%. In contrast, both oil strategies generate relatively high average expected returns of 10.99% and 19.54%. Thereby, the SR of Strategy 2 exceeds one, which is considered as acceptable to good by investors, thus can be regarded as the best trading strategy.

Fourth, similarly to the case of Canada, the oil strategies generate higher returns compared to a buy-and-hold strategy on the Norwegian stock index. Strategy 1 and Strategy 2 produce an expected return of 17.00% and 31.64%, respectively. Both Sharpe ratios are greater than one, thus can be considered as a good investment strategy. Again, Strategy 1 records the lowest volatility, but Strategy 2 represents the best investment in terms of risk-return trade-off.

Finally, in the case of Brazil, the three strategies move in proximity up until January 2020. Here, the Brazilian stock market records a sudden decrease in value associated with the crisis caused by the Covid-19 pandemic. In the same month, the returns of both Strategy 1 and Strategy 2 increase, due to investments in the risk-free rate and shorting, respectively. Altogether, Strategy 2 generates a surprisingly high expected return of 36.12%. This can be attributed to the high volatility of the stock index (36.80%), which often provides opportunity for Strategy 2 to short and thus obtain returns in both up and down states of the market.

To conclude, both oil strategies deliver a better performance than a simple buy-and-hold strategy on the country's stock index, in terms of both returns and Sharpe ratios. For all the five countries, Strategy 2 generates the highest average annualized returns and the highest Sharpe ratios. This can be attributed to the fact that Strategy 2 includes the option to short, and therefore can profit in both up and down turns of the market. This can be particularly seen in the year 2020, where the stock market plunged during the Covid-19 recession. Strategy 1 exhibits the lowest volatility, which can be attributed to the option of investing the short-term T-Bills in case of low expected returns. This possibility for diversification and risk reduction makes Strategy 1 a relatively stable investment opportunity.

Despite the good performance of both oil strategies, we have to interpret these results with caution. The used model has relatively low predictive power, as it is basically impossible to accurately predict stock returns based on one variable. However, as this section is simply intended to provide real-life implications of our research hypotheses, the developed trading strategies astrain from any further level of sophistication. For those interested in more accurately forecasting the stock returns, we recommend including more predictors and using more appropriate forecasting methods such as machine learning models. Despite the simplicity of our implementation, the evidence of our developed trading strategies provides interesting insights for investors, as it points towards the idea that there lies predictable power in oil prices.

Table 13: Mean, Standard Deviation and Sharpe Ratio of the Trading Strategies

	Mean (yearly in %)	SD (yearly in %)	Sharpe ratio (SR)
		Buy-and-hold strategy	
Mexico	-1.72%	23.35%	-0.1089
Hong Kong	8.05%	16.43%	0.4205
Canada	2.46%	16.35%	0.1111
Norway	1.71%	21.36%	0.0419
Brazil	10.15%	36.80%	0.2398
		Strategy 1	
Mexico	8.56%	12.34%	0.6267
Hong Kong	11.65%	12.57%	0.8335
Canada	10.99%	11.51%	0.8966
Norway	17.00%	15.32%	1.0548
Brazil	24.16%	24.85%	0.9169
		Strategy 2	
Mexico	18.07%	22.77%	0.7577
Hong Kong	13.52%	16.13%	0.7679
Canada	19.54%	15.35%	1.2276
Norway	31.64%	19.32%	1.5943
Brazil	36.12%	35.41%	0.9806

Figure 7: Return of the Trading Strategies

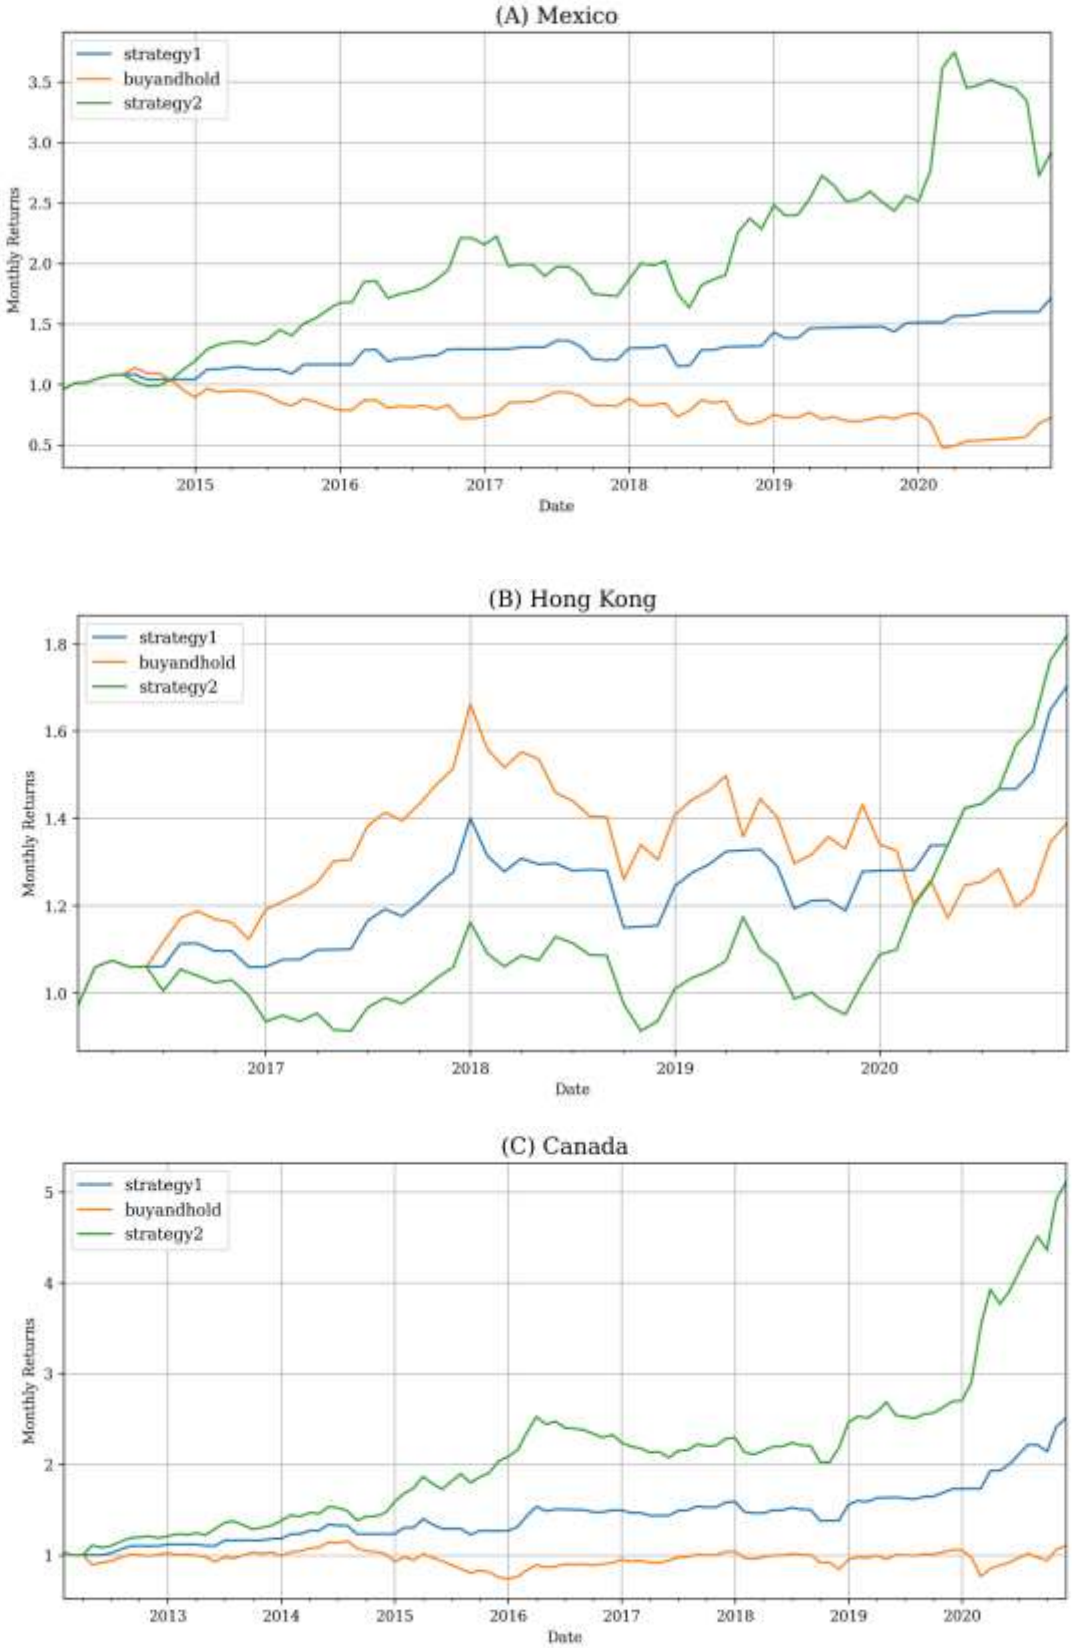

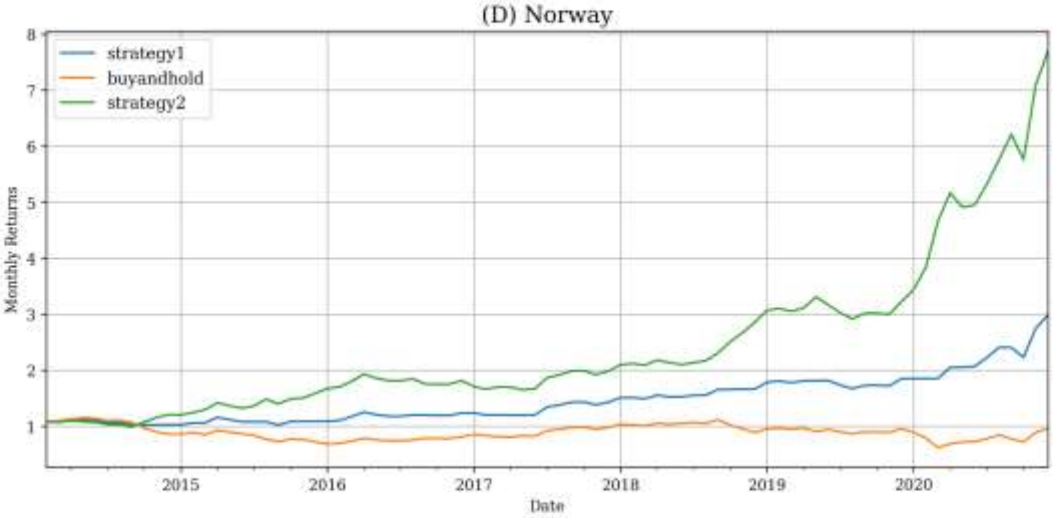
(D) Norway
strategy1
buyandhold
strategy2
Monthly Returns
Date
2015
2016
2017
2018
2019
2020

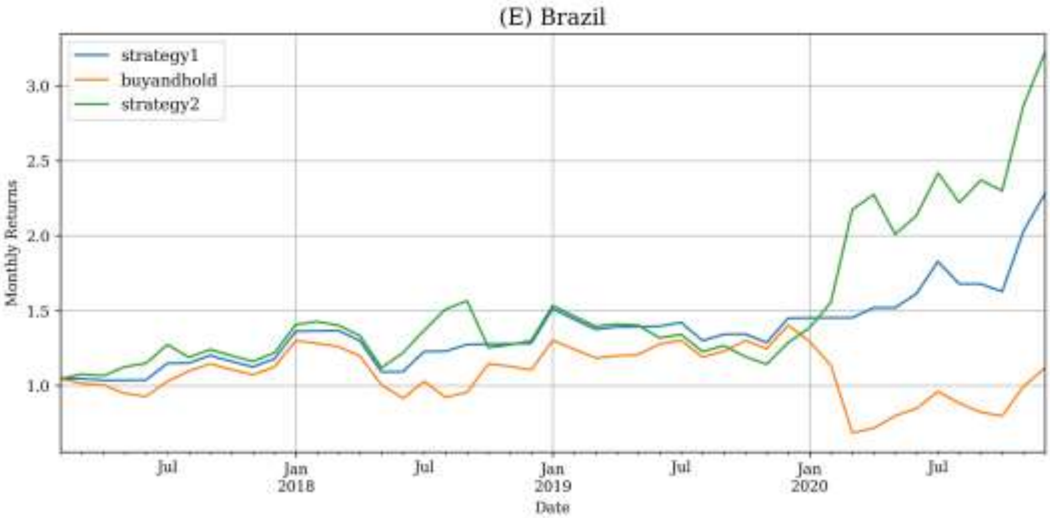
(E) Brazil
strategy1
buyandhold
strategy2
Monthly Returns
Date
Jul
Jan
2018
Jan
2019
Jan
2020

11. Conclusion

The aim of this paper is to research the relationship between oil prices and stock market returns. The paper is divided into two main research hypotheses. The first hypothesis tests whether oil price changes have a contemporaneous impact on the stock market returns. The second hypothesis tests for investors' underreaction to news about oil prices. Both hypotheses are tested on six primarily oil-exporting countries (Canada, Russia, Norway, Brazil, Saudi Arabia and Mexico) and four primarily oil-importing countries (United States, China, Jordan and Hong Kong) over the period March 1983 to December 2020.

The results of the first hypothesis indicate a heterogeneous response of the different countries to oil price changes. We find a significant relationship for all countries at least at some point in time, confirming the existence of a contemporaneous relationship. The results of the analysis provide some interesting observations, namely that oil-exporting countries generally show a positive oil-stock market relationship. Only for oil-exporter Saudi Arabia this not per se holds, as the relationship is extremely volatile and therefore shows little significance. The oil-importing countries show less homogeneous results. Expected would be a negative relationship, however, some countries show contradicting results during specific periods in time, which can possibly be attributed to endogenous factors.

After the oil-stock market relationship is established, we test our second hypothesis, to see whether investors react with a delay to information about oil price changes. The obtained results show that there is an increase in significance and explanatory power for the majority of the countries when adding lags of several trading days. Based on these results, we conclude that there is a delayed reaction of investors. A possible explanation is that investors react at different points in time to changes in oil prices, or have difficulty in assessing the exact impact of these changes on the value of stocks, which is consistent with the gradual diffusion hypothesis proposed by Hong and Stein (1999).

Finally, to demonstrate the implications and applicability of our findings, we develop two trading strategies based on the oil effect, while taking investor underreaction into consideration. Two oil strategies are back tested across five countries and compared to a simple buy-and-hold strategy. Both oil strategies perform better in terms of risk and return as opposed to the buy-and-hold strategy. However, as the strategies astrain from any level of sophistication, further research should consider other variables to include in the model to forecast stock returns more accurately and provide interesting implications for investors.

The limitations of this study are primarily related to the methodology, as it is not able to account for quick changes in coefficient and is sensitive to the window size. Further, due to a probable endogeneity problem, the results might be biased to some extent. Despite these shortcomings, we believe our work adds on the existing literature and is a good basis for understanding the complicated relationship between oil prices and stock markets. More precisely, the findings of the first hypothesis provide a general understanding about the directionality of this time-varying relationship. Thereby, the evidence of our developed trading strategies based on the research hypotheses provides interesting insights for investors, as it points towards the idea that there lies predictable power in oil prices.

.

Bibliography

Alekhina, V., & Yoshino, N. (2018). *Impact of world oil prices on an energy exporting economy incuding monetary policy* (No. 828; ADBI Working Paper Series).

Al-Fayoumi, N. (2009). Oil Prices and Stock Market Returns in Oil Importing Countries: The Case of Turkey, Tunisia and Jordan. *European Journal of Economics, Finance and Administrative Sciences*, *16*(16), 84–98.

Almeida, V., & Pereira Câmara Leal, R. (2015). A joint experimental analysis of investor behavior in ipo pricing methods. *Revista de Administração de Empresas*, *55*(1), 14–25.

Alquist, R., & Kilian, L. (2010). What do we learn from the price of crude oil futures? *Journal of Applied Econometrics*, *25*(4), 539–573.

Anatolyev, S., & Kitov, V. (2007). Using All Observations when Forecasting under Structural Breaks. *Finnish Economic Papers*, *20*(2), 166–176.

Antonakakis, N., Chatziantoniou, I., & Filis, G. (2013). Dynamic co-movements of stock market returns, implied volatility and policy uncertainty. *Economics Letters*, *120*(1), 87–92.

Apergis, N., & Miller, S. M. (2009). Do structural oil-market shocks affect stock prices? *Energy Economics*, *31*(4), 569–575.

Arouri, M., & Nguyen, D. K. (2010). Oil prices, stock markets and portfolio investment: Evidence from sector analysis in Europe over the last decade. *Energy Policy*, *38*(8), 4528–4539.

Asteriou, D., & 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*(C), 204–211.

Ayadi, O. F. (2005). Oil price fluctuations and the Nigerian economy. *OPEC Review*, *29*(3), 199–217.

Bahr, R., & Nikolova, B. (2010). Global oil prices, oil industry and equity returns: Russian experience. *Scottish Journal of Political Economy*, *57*(2), 169–186.

Balvers, R. J., Cosimano, T. F., & McDonald, B. (1990). Predicting Stock Returns in an Efficient Market. *The Journal of Finance*.

Barberis, N., Shleifer, A., & Vishny, R. (1998). A model of investor sentiment. *Journal of Financial Economics*, *49*, 307–343.

Barberis, N., & Thaler, R. (2002). A Survey of Behavioral Finance. In *NBER Working Paper* (No. 9222; NBER Working Papers).

Barsky, R. B., & Kilian, L. (2001). *Do We Really Know that Oil Caused the Great Stagflation? A Monetary Alternative* (No. 8389).

Bernanke, B. S. (1983). *Non-Monetary Effects of the Financial Crisis in the Propagation of the Great Depression* (Working Paper Series 1054). https://www.nber.org/system/files/working_papers/w1054/w1054.pdf

Berument, H. M., Ceylan, N. B., & Dogan, N. (2010). The Impact of Oil Price Shocks on the Economic Growth of Selected MENA Countries. *The Energy Journal*, *31*(1), 149–176.

Broadstock, D., & Filis, G. (2014). Oil price shocks and stock market returns: New evidence from the United States and China. *Journal of International Financial Markets, Institutions and Money*, *33*(C), 417–433.

Brooks, Chris. (2008). *Introductory Econometrics for Finance* (Cambridge University Press, Ed.; Second Edi).

Brown, S., & Yucel, M. (2002). Energy prices and aggregate economic activity: an interpretative survey. *The Quarterly Review of Economics and Finance*, *42*(2), 193–208. https://econpapers.repec.org/RePEc:eee:quaeco:v:%0A42:y:2002:i:2:p:193-208.%0A

Burbidge, J., & Harrison, A. (1984). Testing for the Effects of Oil-Price Rises using Vector Autoregressions. *International Economic Review*, *25*(2), 459–484.

Büyükşahin, B., & Robe, M. (2011). *Does "Paper Oil" Matter? Energy Markets' Financialization and Equity-Commodity Co-Movements* (Bank of Canada Working Paper).

Cai, Z. (2007). Trending time-varying coefficient time series models with serially correlated errors. *Journal of Econometrics*, *136*, 163–188.

Camp, K. M., Mead, D., Reed, S. B., Sitter, C., & Wasilewski, D. (2020). *From the barrel to the pump: the impact of the COVID-19 pandemic on prices for petroleum products*. Monthly Labor Review, U.S. Bureau of Labor Statistics. https://doi.org/10.21916/mlr.2020.24

Campbell, J.Y, Lo, A.W and MacKinlay, A. C. (1997). *The Econometrics of Financial Markets*.

Caporin, M., Chang, C. L., & McAleer, M. (2019). Are the S&P 500 index and crude oil, natural gas and ethanol futures related for intra-day data? *International Review of Economics and Finance*, *59*, 50–70.

Chen, N-F., Roll, R., Ross, S. (1986). Economic forces and the stock market. *The Journal of Business*, *59*(3), 383–403. https://doi.org/10.1016/j.jempfin.2005.09.001

Chen, T. (2013). Do Investors Herd in Global Stock Markets? *Journal of Behavioral Finance*, *14*(3), 230–239.

Cheng, I.-H., & Xiong, W. (2013). *The Financialization of Commodity Markets* (No. 19642; NBER Working Paper Series).

Ciner, C. (2001). Energy Shocks and Financial Markets: Nonlinear Linkages. *Studies in Nonlinear Dynamics & Econometrics*, *5*(3), 1–11.

Cologni, A., & Manera, M. (2008). Oil prices, inflation and interest rates in a structural cointegrated VAR model for the G-7 countries. *Energy Economics*, *30*(3), 856–888. https://econpapers.repec.org/RePEc:eee:%0Aeneeco:v:30:y:2008:i:3:p:856-888.%0A

Cong, R. G., Wei, Y. M., Jiao, J. L., & Fan, Y. (2008). Relationships between oil price shocks and stock market: An empirical analysis from China. *Energy Policy*, *36*(9), 3544–3553. https://doi.org/10.1016/j.enpol.2008.06.006

Cuthbertson, K., & Nitzsche, D. (2004). *Quantitative Financial Economics: Stocks, Bonds and Foreign Exchange* (John Wiley & Sons, Ed.; 2nd Editio).

Darrat, A. F., Zhong, M., & Cheng, L. T. W. (2007). Intraday volume and volatility relations with and without public news. *Journal of Banking and Finance*, *31*(9), 2711–2729. https://www.sciencedirect.com/science/article/abs/pii/S0378426607000568?via%3Dihub

DeLong, J. B., Schleifer, A., & Summers, Lawrence Waldman, R. (n.d.). *Positive Feedback Investment Strategies and Destabilizing Rational Specul*.

Driesprong, G., Jacobsen, B., & Maat, B. (2008). Striking oil: Another puzzle? *Journal of Financial Economics*, *89*(2), 307–327. https://doi.org/10.1016/j.jfineco.2007.07.008

Edelstein, P., & Kilian, L. (2009). How sensitive are consumer expenditures to retail energy prices? *Journal of Monerary Economics*, *56*(6), 766–779. https://ideas.repec.org/a/eee/moneco/v56y2009i6p766-779.html.

El-Sharif, I., Brown, D., Burton, B., Nixon, B., & Russell, A. (2005). Evidence on the nature and extent of the relationship between oil prices and equity values in the UK. *Energy Economics*, *27*(6), 819–830.

Fama, E. F. (1970). Efficient Capital Markets : A Review of Theory and Empirical Work. *The Journal of Finance*, *25*(2), 383–417.

Fan, Q., & Jahan-Parvar, M. R. (2011). U.S. Industry-Level Returns and Oil Prices. *International Review of Economics and Finance*, *22*(1), 112–128. https://www.sciencedirect.com/science/article/pii/S1059056011001109?casa_token=_VQz6RlY4MoAAAAA:A6u7duzCXLpp7wJ9RxVzajIoizq-4jpqcusy9cQfkKmMJNXk87g6YUhlvAVa_bqK7t-Allz4

Federal Reserve Bank of St.Louis. (n.d.). *Industrial Production: Total Index [INDPRO]*. https://fred.stlouisfed.org/series/INDPRO

Frazzini, A. (2006). The Disposition Effect and Underreaction to News. *The Journal of Finance*, *61*(4), 2017–2046. https://onlinelibrary.wiley.com/doi/10.1111/j.1540-6261.2006.00896.x

George, T. J., & Hwang, C.-Y. (2005). The 52-Week high and momentum investing. *The Journal of Finance*, *59*(5), 2145–2176. https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1540-6261.2004.00695.x

Gisser, M., & Goodwin, T. H. (1986). Crude Oil and the Macroeconomy : Tests of Some Popular Notions : Note. *Journal of Money, Credit and Banking*, *18*(1), 95–103.

Gourène, G. A. Z., & Mendy, P. (2018). Oil prices and African stock markets co-movement: A time and frequency analysis. *Journal of African Trade*, *5*(1–2), 55–67. https://doi.org/10.1016/j.joat.2018.03.002

Government of Canada. (2020). *Crude oil facts*. https://www.nrcan.gc.ca/science-data/data-analysis/energy-data-analysis/energy-facts/crude-oil-facts/20064

Grossman, S. J., & Stiglitz, J. E. (1980). On the Impossibility of Informationally Efficient Markets. *The American Economic Review*, *70*(3), 393–408.

Hamilton, J. D. (1983). Oil and the macroeconomy since world war II. *Journal of Political Economy*, *91*(2), 228–248. https://doi.org/10.1086/261140

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

Hamilton, J. D. (2013). Oil prices, exhaustible resources and economic growth. *Handbook on Energy and Climate Change*, *1*, 29–63.

Hong, H., & Stein, J. C. (1999). A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets. *The 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. https://doi.org/10.1016/j.jfineco.2005.09.010

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

Imbert, F. (2020). *Dow drops nearly 3,000 points, as coronavirus collapse continues; worst day since '87*. CNBC. https://www.cnbc.com/2020/03/15/traders-await-futures-open-after-fed-cuts-rates-launches-easing-program.html

Inoue, A., & Rossi, B. (2012). Out-of-Sample Forecast Tests Robust to the Choice of Window Size. *Journal of Business & Economic Statistics*, *30*(3), 432–453.

Jones, C. M., & Kaul, G. (1996). Oil and the Stock Markets. In *The Journal of Finance* (Vol. 51, Issue 2).

Kahneman, D., & Tversky, A. (1974). Judgment under Uncertainty: Heuristics and Biases. *Sciences, New Series*, *185*(4157), 1124–1131.

Kilian, L. (n.d.). Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy? *The Review of Economics and Statistics*, *90*(2), 216–240.

Kilian, L., & Park, C. (2009). The Impact of Oil Price Shocks on the U.S. Stock Market. *International Economic Review*, *50*(4), 1267–1287.

Lamont, O. A., & Thaler, R. H. (2003). Anomalies: The law of one price in financial markets. *Journal of Economic Perspectives*, *17*(4), 191–202. https://doi.org/10.1257/089533003772034952

Latif, M., Arshad, S., Fatima, M., & Farooq, S. (2011). Market efficiency, market anomalies, causes, evidences, and some behavioral aspects of market anomalies. *Research Journal of Finance and Accounting*, *2*(9), 1697–2222.

Lin, C.-C., Fang, C., & Cheng, H.-P. (2010). Relationships between oil price shocks and stock market: an empirical analysis from Greater China. *China Economic Journal*, *3*(3), 241–254.

Locke, P. R., & Mann, S. C. (2000). *Do Professional Traders Exhibit Loss Realization Aversion*. https://doi.org/10.2139/ssrn.251942

Loungani, P. (1986). Oil Price Shocks and the Dispersion Hypothesis. *The Review of Economics and Statistics*, *68*(3), 536–539.

Maghyereh, A. (2004). Oil Price Shocks and Emerging Stock Markets: A Generalized VAR Approach. *International Journal of Applied Econometrics and Quantitative Studies*, *1*(2), 27–40.

Merton, R. C. (1987). A Simple Model of Capital Market Equilibrium with Incomplete Information. *The Journal of Finance*, *42*(3), 483–510. https://doi.org/10.1111/j.1540-6261.1987.tb04565.x

Miller, J. I., & Ratti, R. A. (2009). Crude oil and stock markets: Stability, instability, and bubbles. *Energy Economics*, *31*(4), 559–568. https://doi.org/10.1016/j.eneco.2009.01.009

Mohanty, S. K., & Nandha, M. (2011). Oil Risk Exposure: The Case of the U.S. Oil and Gas Sector. *The Financial Review*, *46*(1), 165–191. https:%0A//ideas.repec.org/a/bla/finrev/v46y2011i1p165-191.html.%0A

Munk, C. (2019). *Financial Markets and Investments: Notes*. Copenhagen Business School.

Nandha, M., & Faff, R. (2008). Does oil move equity prices? A global view. *Energy Economics*, *30*(3), 986–997.

Narayan, P., & Narayan, S. (2010). Modelling the impact of oil prices on Vietnam's stock prices. *Applied Energy*, *87*(1), 356–361.

Nguyen, C., & Bhatti, M. (2012). Copula model dependency between oil prices and stock markets: Evidence from China and Vietnam. *Journal of International Financial Markets*, *22*(4), 758–733.

Nickerson, R. S. (1998). Confirmation Bias: A Ubiquitous Phenomenon in Many Guises. *Review of General Psychology*, *2*(2), 175–220.

Norsk Petroleum. (2020). *Exports of Oil and Gas*. https://www.norskpetroleum.no/en/production-and-exports/exports-of-oil-and-gas/

OPEC. (n.d.). *Member Countries*. https://www.opec.org/opec_web/en/about_us/25.htm#:~:text=The Organization of the Petroleum,Founder Members of the Organization.

OPEC. (2019). *Saudi Arabia facts and figures*. https://www.opec.org/opec_web/en/about_us/169.htm

Papapetrou, E. (2001). Oil price shocks, stock market, economic activity and employment in Greece. *Energy Economics*, *23*(5), 511–532.

Park, J., & Ratt. (2008). Oil price shocks and stock markets in the U.S. and 13 European countries. *Energy Economics*, *30*(5), 2587–2608.

Pedersen, L. H. (2015). *Efficiently Inefficient: How Smart Money Invests and Market Prices Are Determined*. Priceton University Press.

Pesaran, M., & Timmermann, A. (2007). Selection of estimation window in the presence of breaks. *Journal of Econometrics*, *137*(1), 134–161.

Plous, S. (1993). *The psychology of judgment and decision making* (Mcgraw-Hill Book Company, Ed.).

Pollet, J. M. (2011). *Predicting Asset Returns With Expected Oil Price Changes*. https://doi.org/10.2139/ssrn.722201

Pompian, M. (2006). *Behavioral Finance and Wealth Management: How to Build Optimal Portfolios That Account for Investor Biases* (Inc. John Wiley & Sons, Ed.).

Pouget, S., & Sauvagnat, Julien Villeneuve, S. (2017). A Mind Is a Terrible Thing to Change: Confirmatory Bias in Financial Markets. *Review of Financial Studies*, *30*(6), 2066–2109.

Sadorsky, P. (34 C.E.). Correlations and volatility spillovers between oil prices and the stock prices of clean energy and technology companies. *2012*, *1*(248–255).

Sadorsky, P. (1999). Oil price shocks and stock market activity. *Energy Economics*, *21*(5), 449–469.

Sadorsky, P. (2001). Risk factors in stock returns of Canadian oil and gas companies. *Energy Economics*, *23*(1), 17–28.

Schwert, G. W. (1990). Stock Returns and Real Activity: A Century of Evidence. *The Journal of Finance*, *45*(4), 1237–1257.

Shapiro, M. D. (1989). *Assessing the Federal Reserve's Measures of Capacity and Utilization* (Brookings Papers on Economic Activity).

Shefrin, H., & Statman, M. (1985). The Disposition to Sell Winners Too Early and Ride Losers Too Long: Theory and Evidence. *The Journal of Finance*, *40*(3), 777–790. https://doi.org/10.1111/j.1540-6261.1985.tb05002.x

Shiller, R. J. (1999). Human behavior and the efficiency of the financial system. *Handbook of Macroeconomics*, *1*(Chapter 20), 1305–1340. https://doi.org/10.1016/S1574-0048(99)10033-8

Tobben, S. (2020). *Oil for Less Than Nothing? Here's How That Happened*. Bloomberg.

U.S. Energy Information Administration. (2017). *Russia*. https://www.eia.gov/international/analysis/country/RUS

U.S. Energy Information Administration. (2020a). *China - Overview*. https://www.eia.gov/international/analysis/country/CHN

U.S. Energy Information Administration. (2020b). *Despite the U.S. becoming a net petroleum exporter, most regions are still net importers*. https://www.eia.gov/todayinenergy/detail.php?id=42735#:~:text=In September 2019%2C the United,b%2Fd of crude oil

Wang, Y., Wu, C., & Yang, L. (2013). Oil price shocks and stock market activities: Evidence from oil-importing and oil-exporting countries. *Journal of Comparative Economics*, *41*(4), 1220–1239.

Williams, J. B. (1938). *The Theory of Investment Value*. Cambridge: Harvard University Press.

Wooldridge, J. M. (2008). *Introductory Econometrics: A Modern Approach* (C. Learning, Ed.).

Youssef, M., & Mokni, K. (2019). Do crude oil prices drive the relationship between stock markets of oil-importing and oil-exporting countries? *Economies*, *7*(3), 1–22. https://doi.org/10.3390/economies7030070

Zhang, Y. J., & Wang, Z. Y. (2013). Investigating the price discovery and risk transfer functions in the crude oil and gasoline futures markets: some empirical evidence. *Applied Energy*, *104*(1), 220–228.

Appendix A

Table A1: Correlation Matrix of the Stock Indices

	United States	Mexico	Canada	China	Norway	Saudi Arabia	Brazil	Russia	Jordan	Hong Kong
United States	1.0000	0.5188	0.7895	0.1678	0.6847	0.3024	0.3924	0.4839	0.1981	0.5822
Mexico	0.5188	1.0000	0.5299	0.0892	0.4917	0.2039	0.4327	0.6112	0.1411	0.4373
Canada	0.7895	0.5299	1.0000	0.1749	0.7132	0.3461	0.4155	0.5610	0.2139	0.5988
China	0.1678	0.0892	0.1749	1.0000	0.1435	0.1527	0.1230	0.2214	-0.0505	0.2322
Norway	0.6847	0.4917	0.7132	0.1435	1.0000	0.3791	0.3968	0.5148	0.2434	0.5737
Saudi Arabia	0.3024	0.2039	0.3461	0.1527	0.3791	1.0000	0.1959	0.2430	0.3527	0.2673
Brazil	0.3924	0.4327	0.4155	0.1230	0.3968	0.1959	1.0000	0.6025	0.1423	0.3854
Russia	0.4839	0.6112	0.5610	0.2214	0.5148	0.2430	0.6025	1.0000	0.1304	0.5204
Jordan	0.1981	0.1411	0.2139	-0.0505	0.2434	0.3527	0.1423	0.1304	1.0000	0.1686
Hong Kong	0.5822	0.4373	0.5988	0.2322	0.5737	0.2673	0.3854	0.5204	0.1686	1.0000

Appendix B

Figure B1: Autocorrelation Plot WTI Futures

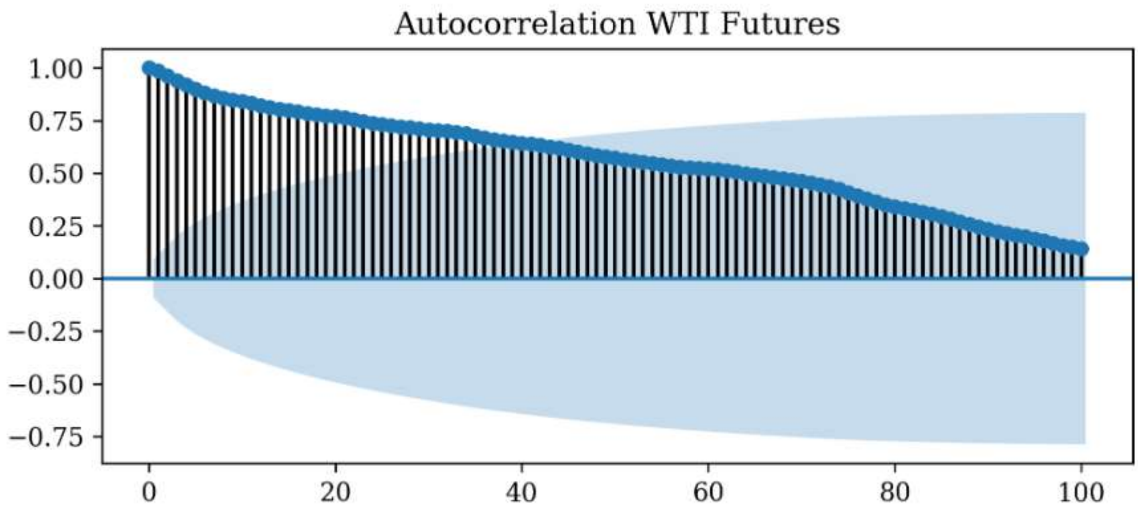

Figure B2: Autocorrelation Plot Industrial Production Index

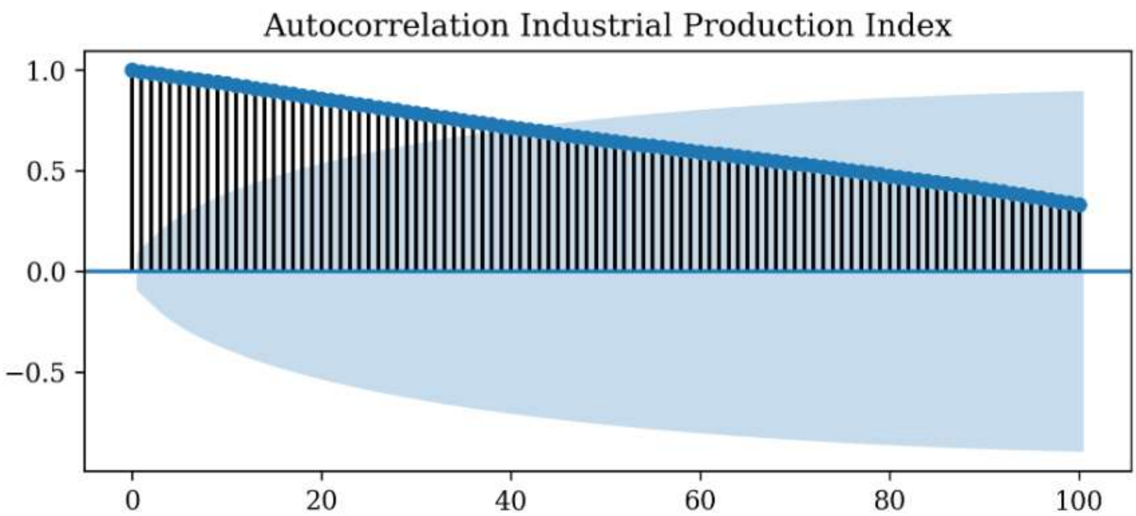

Figure B3: Autocorrelation Plot Stock Market Indices

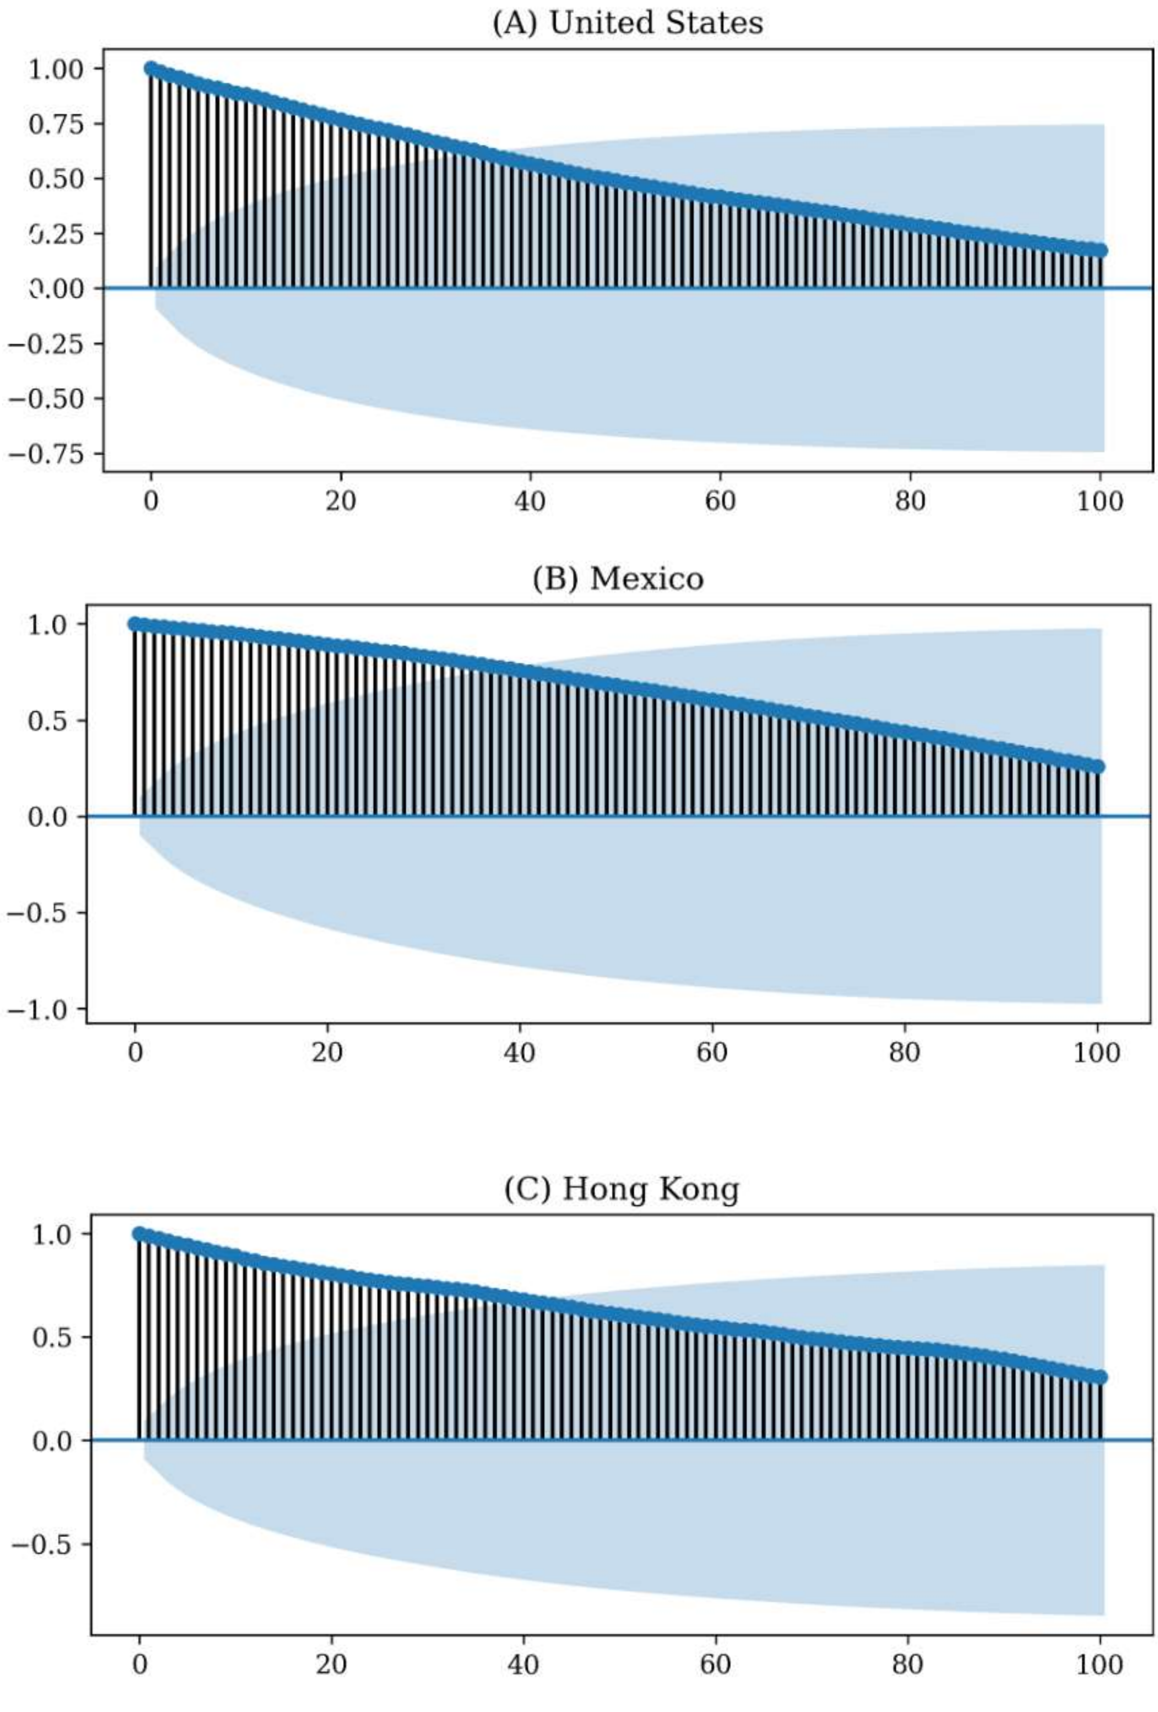

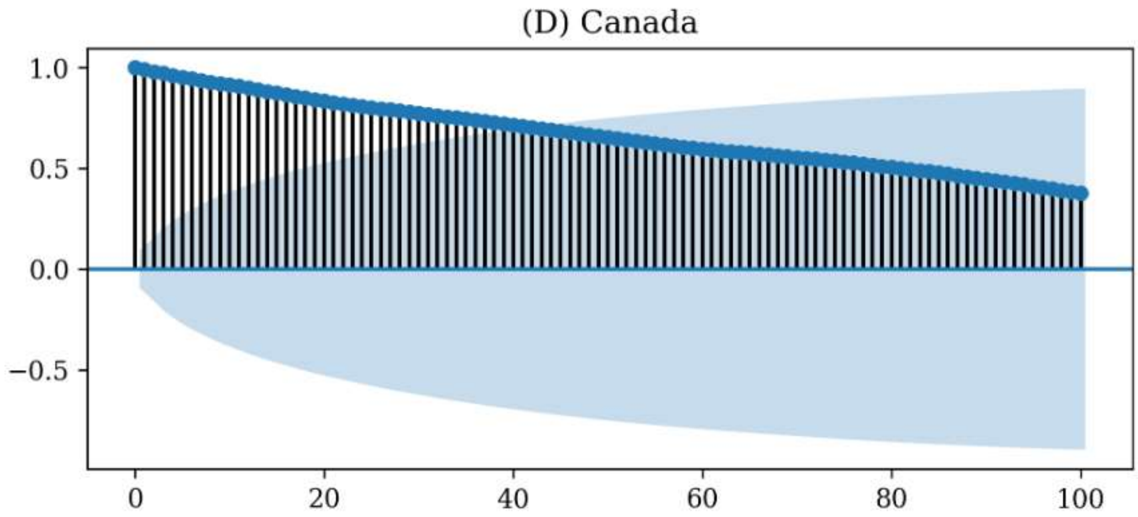
(D) Canada
1.0
0.5
0.0
−0.5
0
20
40
60
80
100

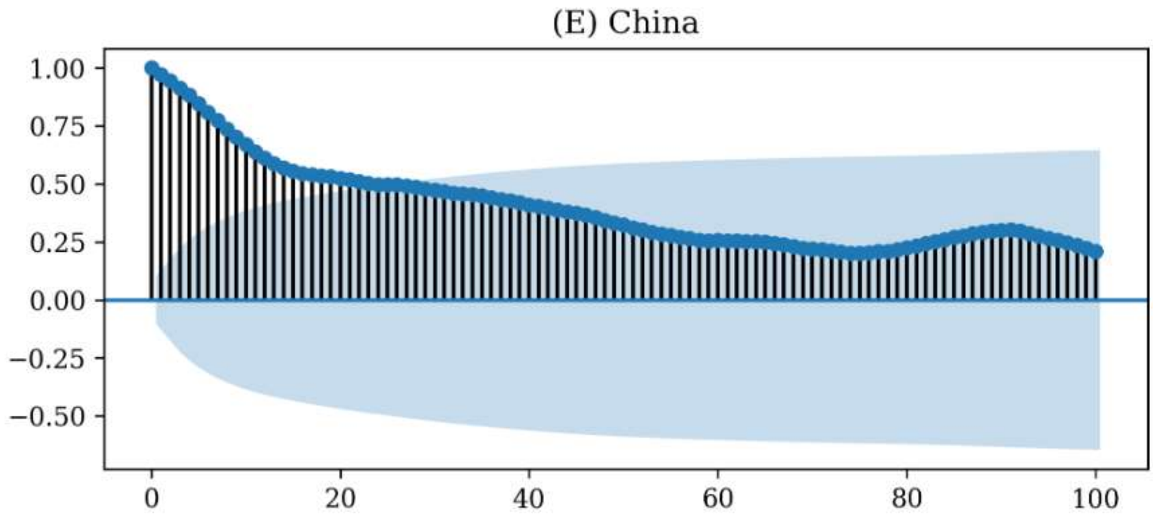
(E) China
1.00
0.75
0.50
0.25
0.00
−0.25
−0.50
0
20
40
60
80
100

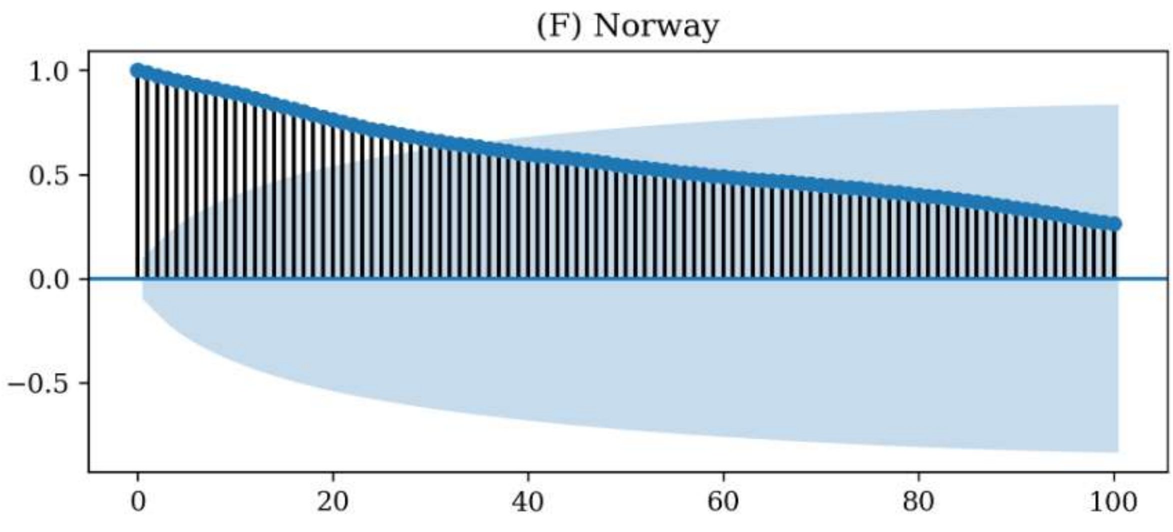
(F) Norway
1.0
0.5
0.0
−0.5
0
20
40
60
80
100

(G) Jordan

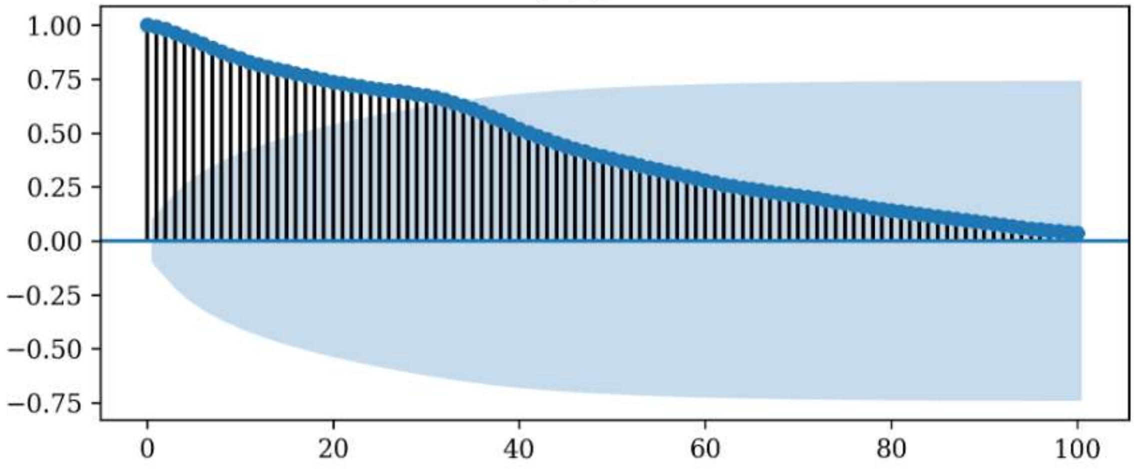

(H) Saudi Arabia

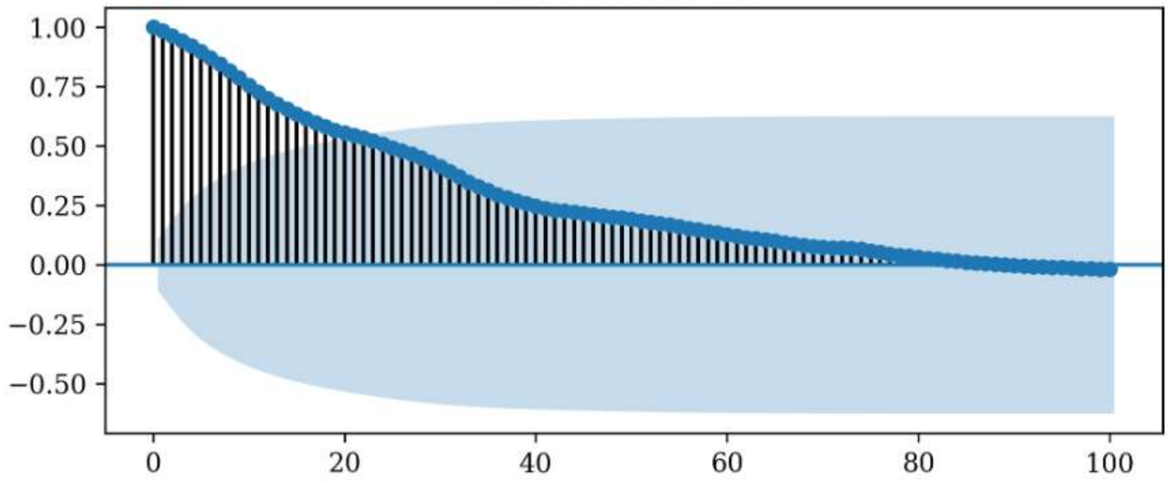

(I) Brazil

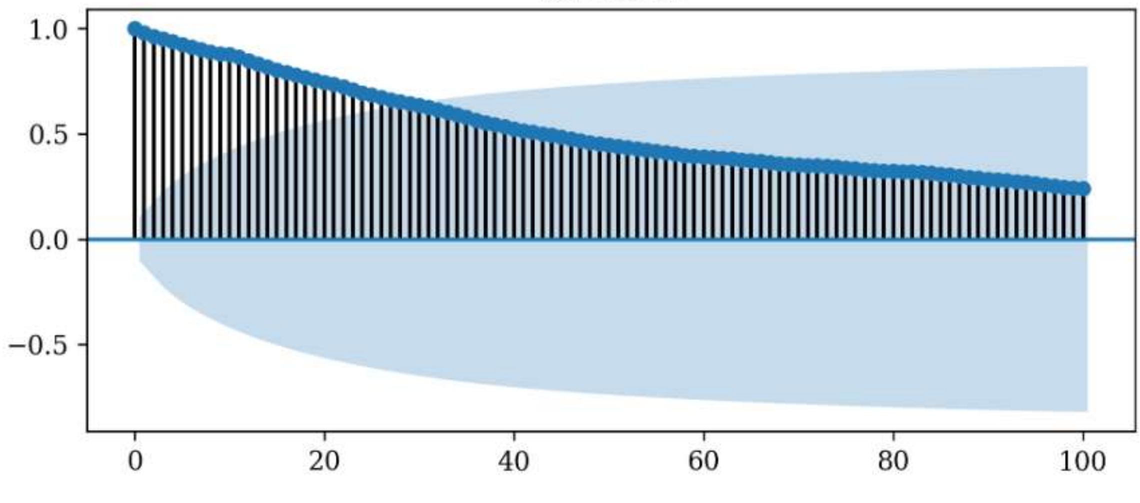

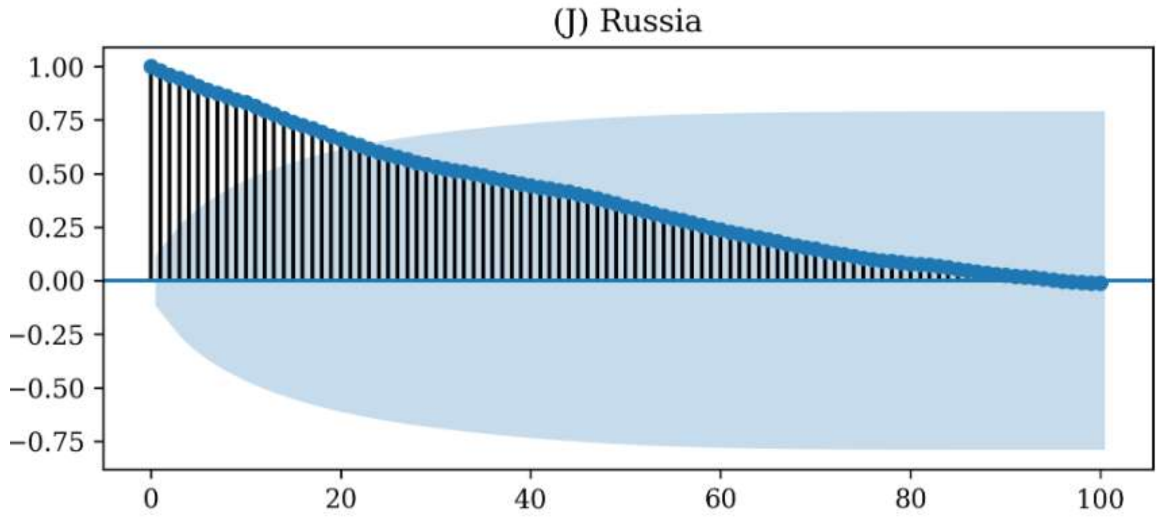

(J) Russia
1.00
0.75
0.50
0.25
0.00
−0.25
−0.50
−0.75
0
20
40
60
80
100