Johan L. S. Karlsen
Morten B. Mjøen

# Forecasting the Norwegian Krone Exchange Rate using the Oil Price: A Trader's and a Statistician's Perspective

Master's thesis in Financial economics
Trondheim, May 2015

Norwegian University of Science and Technology
Faculty of Social Sciences and Technology Management
Department of Economics

NTNU
Kunnskap for en bedre verden

## Preface

We would like to thank our supervisor, Professor Xunhua Su, for superior guidance, excellent ideas and enjoyable collaboration. His availability to have meetings and discussions with us is highly appreciated and has very much contributed to the making as well as the final result of this thesis. Further, we would like to thank Åge Halvorsen and Kristoffer Dahlberg from Sparebank 1 Markets for providing us with valuable data sets. We would also like to thank the other Professors on the faculty and fellow students for making the last two years so exciting.

Trondheim 8. May 2015

Sincerely

Johan L. S. Karlsen
Morten B. Mjøen

## Abstract

In the past nearly one year, there has been a 30% depreciation of the Norwegian currency. The dramatic fall in the oil price is blamed to be the key reason since the oil price drop can be considered as an exogenous shock to the oil dependent Norwegian economy. If this is true, can the oil price predict movements in the Norwegian krone? We examine this issue from two different perspectives. First, we use high frequency data such as daily and hourly to simulate simple strategies involving blindly trading the dollar based on signals given by the oil price. We find that when using the direction of change in the oil price as a predictor for the direction of change in the USD/NOK exchange rate we are able earn higher risk-adjusted returns than a simple buy and hold strategy. Second, in spirit of Ferraro et al. (2015), we try to forecast daily and hourly changes in the exchange rate using oil price changes as the only predictor. We find that contemporaneous changes in the oil price significantly outperforms the random walk in terms of forecasting ability, while lagged changes in the oil price yield indistinct results highly dependent on the timing of the exchange rate data.

## Table of Contents

# Figures



# Tables

# 1 Introduction

Between the 19th of June 2014 and the 13th of January 2015 the US dollar appreciated 26% against the Norwegian krone while the Brent oil spot price plummeted 61%. "*There is no doubt that the latest movements in the exchange rate are related to the decrease in the oil price*" [1], said the Governor of the Norwegian Central Bank, Øystein Olsen, in the middle of this financial turmoil. At that day the USD/NOK exchange rate and the Brent oil spot price was 6,60 and 86,38 USD/barrel respectively. At the 13th of January 2015, less than three months later, a barrel of Brent oil was priced 48% lower and the US dollar had appreciated another 17% against the Norwegian krone. One could say that this is a state of crisis for the Norwegian oil sector while a state of euphoria for other exporting businesses in Norway. We have no reason to believe that the state of euphoria outweigh the state of crisis knowing how important the petroleum sector is for the Norwegian economy.[2] Figure 1 plots the daily time series of the Brent oil spot price and the Norwegian exchange rate expressed as how many US dollars one Norwegian krone can buy. It speaks in favor of the belief of the Norwegian Central Bank. The figure illustrates how closely the two variables have moved since the early 2000s.[3] It certainly suggests the possibility of a causal relationship between the oil price and the exchange rate of Norway. The oil price is a leading economic and financial variable that is commonly referred to as an important driver of the world economy (Ghalayaini 2011). Moreover, the price of oil is denominated in dollars and is traded on highly centralized financial markets. The Norwegian krone, on the other hand, is a currency that is relatively insignificant in the global currency market and can't be expected to determine the state of the global economy or the price of oil. It seems, however, more likely that the price of oil has an impact on the exchange rate of Norway. Jarle Bergo (2004), a former Vice Governor of the Norwegian Central Bank, mentions the tendency that the Norwegian krone appreciates when production in Norway is high, and depreciates when production is low. There is no doubt that the petroleum sector is a major part of production in Norway and that the oil price is critical when valuing this part of total production. The relationship between the oil price and the Norwegian exchange rate proposed above strongly motivates us to analyze whether the oil price can be used to predict movements in the exchange rate.

---

[1] Our translation. Quote from a press conference at the 22nd of October 2014 (Haug 2014).

[2] The petroleum sector accounts for 22% of the Norwegian GDP, 30% of state revenues and about 50% of Norwegian exports (The Norwegian Petroleum Sector 2014).

[3] We have estimated the correlation between the oil price and the exchange rate to be 0,85 based on the period from March 2001 to February 2015.

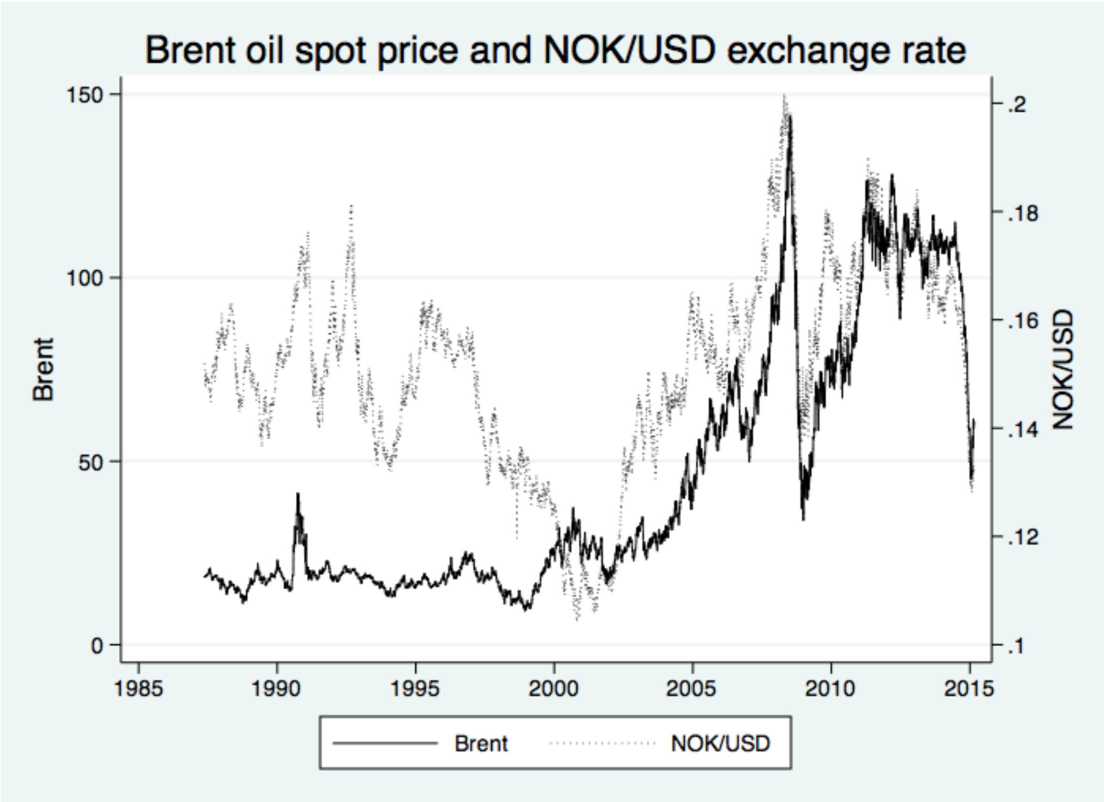


**Figure 1 – Time series of Brent oil spot price and NOK/USD exchange rate**

In this paper we explore the predictive ability of the Brent oil spot price on the Norwegian exchange rate from two different perspectives: A trader's and a statistician's. Their perception of success is very different (Melvin et al. 2013). A paper by Ferraro et al. (2015) named "*Can Oil Prices Forecast Exchange Rates?"* is the foundation for our paper. They present an idea saying that changes in the oil price may have significant explanatory power on the exchange rates of oil abundant countries. Two main models are presented: A *contemporaneous model* and a *true forecasting model*. The former is a simple regression of the first differenced natural logarithm of the USD/NOK exchange rate on the first differenced natural logarithm of the Brent oil spot price. The latter substitutes contemporaneous values of the oil price with lagged values. The results from the two models suggest significant explanatory power of both the oil price and the lagged oil price on the exchange rate; a decrease in the oil price is paired with an increase in the Norwegian exchange rate, vice versa. This relationship motivated us to form different trading strategies based on the idea that oil price changes can predict future exchange rate changes. The idea is simple: If we observe a decrease in the oil price, we long the dollar; if we observe an increase in the oil price, we short the dollar. The strategies are simulated on 4-year daily data and on 140-day hourly data.

We explore trading opportunities by taking only long positions, taking both long and short positions, and by imposing boundaries for which movements in the oil price have to cross before we take any positions. We have imposed the strategies on time-windows of different lengths as well as the whole sample. The windows are rolled through the whole sample to investigate the performance of the strategies over time. Our findings indicate that we are able to construct strategies that outperform a simple buy and hold strategy in terms of risk-adjusted returns for both daily and hourly trading.

In the second part of our analysis we perform a thorough statistical exercise where we investigate the predictive ability of oil price changes on exchange rate changes using the framework of Ferraro et al. (2015). We extend their daily analysis by using data sets recorded at different daily hours, thus exploring the effect of timing, and by analyzing hourly data. However, we omit investigating quarterly and monthly data. Compared to Ferraro et al. (2015) we focus on the Brent oil spot price and the Norwegian krone exchange rate while they mainly focus on the WTI oil spot price and the Canadian dollar exchange rate. To investigate the predictive ability of the oil price on the exchange rate we first estimate one-step-ahead out-of-sample forecasts for the exchange rate. We do this by using different rolling in-sample estimation windows for both the contemporaneous model and the true forecasting model. For each window size the out-of-sample forecasts are compared to those of a random walk model without drift and evaluated based on the Diebold and Mariano (DM) test statistic (Diebold and Mariano, 1995).

The contemporaneous model uses contemporaneous values of oil price changes to "predict" already realized changes of the exchange rate and is in reality an out-of-sample fit exercise. In practice it is impossible to use such a model to forecast exchange rate changes at a future point in time. However, good performance of such a model documents a strong out-of-sample relationship between the variables. Further, if one has accurate forecasts of future oil price changes this model can prove useful for exchange rate forecasting (Ferraro et al. 2015). When using daily data this model statistically outperforms the random walk for all in-sample window sizes. Using hourly data the model statistically outperforms the random walk for all in-sample window sizes up to and including ½ of the total sample.

The true forecasting model is using lagged oil price changes to predict future changes of the exchange rate and therefore enables us to directly measure forecasting ability. Regarding daily data the results from this exercise are highly dependent on the time of recording of the exchange rate data. Using lagged first differences of the oil price there is actually an information overlap in our daily time series.[4] We identify that the predictive ability of the true forecasting model is better when the information overlap is longer. With an information overlap of three hours, our true forecasting model outperforms the random walk for rolling in-sample window sizes of ¼ of the total sample size and larger. With an information overlap of 30 minutes, the model never outperforms the random walk. We also studied the performance of the true forecasting model over time and found that the model has statistically outperformed the random walk during short periods in the past. For the data with the longest overlap these periods were longer and more frequent.

In chapter 2 we will present previous research we found relevant for our paper. Chapter 3 is a description of the data and of our main variables. In chapter 4 we present various empirical models and ordinary least squares (OLS) regressions to illustrate the statistical relationship between the USD/NOK exchange rate, short-term interest rates and the Brent oil spot price. In chapter 5 we try to utilize the relationship by forming trading strategies where we make trades in the dollar based on signals given by the oil price. Chapter 6 is an investigation of the statistical forecasting ability of the oil price on the exchange rate using the framework presented by Ferraro et al. (2015). Chapter 7 discusses our results in the perspective of a trader and a statistician. Chapter 8 concludes our paper. All regressions in our analysis are ordinary least squares (OLS) and all statistical operations are conducted in Stata 11.

[4] We define *information overlap* as the length of time for which the lagged first differenced oil price overlap the first differenced exchange rate. The overlap is a result of the oil price being recorded *after* the exchange rate each day. I.e. if the oil price is recorded at 16:30 and the exchange rate at 13:30, the overlap is three hours. Further discussed in section 4.2.

# 2 Previous research

A paper by Ferraro et al. (2015) considers data on several commodity prices and exchange rates of several countries and investigates the predictive ability of price changes in a country's major commodity export on its exchange rate. Their analysis involves the price changes of WTI oil, denominated in US dollars, and its predictive ability on the exchange rates of Canada and Norway measured against the dollar. The existence of a short-term relationship between the oil price changes and changes in the nominal exchange rates of the two countries is demonstrated through the use of daily data. First they conduct an out-of-sample fit exercise where a contemporaneous oil price change is used. In this exercise the findings of predictive ability is quite robust. When lagged commodity price change is used predictive ability is also found. The result of the latter exercise appears with less significance and is assorted, meaning that the model only outperforms the random walk in some parts of the sample. The paper is the first to demonstrate, with high statistical significance, short horizon predictive ability of oil prices on exchange rates. Due to the success of the methodology used by Ferraro et al. (2015) it has become the basis for our thesis.

The literature on predicting nominal exchange rates using macroeconomic fundamentals is large and the general view is that traditional theory-based models perform unsatisfactory. Meese and Rogoff (1983) consider a range of exchange rate models and their out-of-sample forecasting ability and find that a random walk model performs as well as any of them. This early paper sheds light on the task of beating the random walk as being the central one to take on and points to the importance of out-of-sample fit when evaluating exchange rate models. Meese (1990) points to research since the 1970s and reports that even models that use contemporaneous values aren't good predictors of exchange rates and that economists do not yet understand the determinants of movements in exchange rates in the short and medium term. He also addresses that short-run behavior of exchange rate market participants can be a challenge for traditional modeling of exchange rates. Mussa (1990) emphasizes that some of the shortcomings of the theory-based models can be attributed to failures and lack of sophistication and technique in analyzing data. Cheung et al. (2005) test several theoretical macroeconomic models developed during the nineties by focusing on out-of-sample prediction ability. The study concludes that none of the models tested are very successful. It seems to be a shared view in the literature that monetary fundamentals haven't been very helpful in forecasting exchange rates (Ferraro et al. 2015, Cheung et al. 2005). Common for the results above is the use of lower frequency data, typically monthly or quarterly. Ferraro et al. (2015) points to some findings of predictive ability of macroeconomic fundamentals but emphasizes that inference procedures have been called into question.

Highly relevant for our study is literature focusing on the ability of commodity prices for explaining fluctuations of exchange rates. In two studies from 1998 Amano and Norden analyze the in-sample relationship between oil prices and exchange rates for the United States, Japan and Germany. They identify a stable long-run relationship between real exchange rates and real oil prices through a dynamic cointegration analysis and an error correction model that provides significantly better forecasts than a random-walk model. More recently Chen and Rogoff (2002) investigate what determines the real exchange rates of Australia, Canada and New Zealand. Findings reveal that prices of commodity exports measured in US dollars appear to have an influence on real exchange rates. This study also emphasizes that since commodity products are traded in highly centralized global markets it can be considered an exogenous source of terms of trade fluctuations. Chen et al. (2010) studied the forecasting abilities of exchange rates and commodity prices in both directions. Their findings reveal that exchange rates are very useful in forecasting out-of-sample commodity prices, but the reverse analysis is not as satisfying. Commodity prices are not found to consistently produce better forecasts for exchange rate movements than a random walk model. Both the 2002 and 2010 studies operate with constructed country-specific commodity price indices, instead of individual commodity prices, and they use quarterly data. A study by the European Central Bank (Habib et al. 2007) uses quarterly data and finds no impact of the real oil price on the real exchange rate of Norway. Regarding the oil price as a variable, ECB (Fratzscher et al. 2014) emphasizes that the oil increasingly has become a financial asset over the last decades and that this “financialization”[5] may be the reason for a closer link between oil prices and other assets, such as equity market returns. The study also points to the rising negative correlation between the US dollar and oil prices and to rising levels and volatility of oil prices.

Although we use the foundation of Ferraro et al. (2015) we omit parts of their analysis in our paper, and take for granted some of the results they find. For instance, they consider a cointegration model proposed by Mark (1995), and find no sign of outperformance of the

[5] Financialization is defined as number of open interest contracts in the oil futures market (Fratzscher et al. 2014)

random walk. They also state that *"(…) imposing cointegration is important at lower frequency data; therefore we expect them not to be important in our analysis on high frequency data*" (Ferraro et al. 2015, p. 28). We further omit all analysis on monthly and quarterly data, since the findings in previous research and, most importantly Ferraro et al. (2015), are not very optimistic regarding low frequency data.

# 3 Data material and descriptive statistics

## 3.1 Data sets

We are using several data sets for the exchange rate that differ in terms of timing, frequency and length. The two main sets both consist of daily recordings of the USD/NOK exchange rate from 20th of May 1987 to 23rd of February 2015. One set is obtained from Bank of England and is recorded at 16:00 UK time each day. The other set is obtained from Norges Bank and is recorded at 13:30 UK time each day. We also make use of several other daily data sets obtained from Bank of England that contain daily recordings at different points in time than the two sets mentioned. These sets contain four years of data and are used to develop a trading strategy and to study the effect of the timing. The 4-year data sets more specifically contain data from the 3rd of May 2010 to the 23rd of February 2015. In addition, we use a set containing 140 days of hourly USD/NOK exchange rate data. The set is obtained from Sparebank1 Markets and contains hourly recordings from 14:00 the 28th of August to 09:00 the 12th of March. We further use a data set on the GBP/NOK exchange rate from the 20th of May 1987 to the 23rd of February 2015 recorded at 16:00 UK time, obtained from Bank of England.

In all daily analysis we use Brent crude oil prices recorded by Thomson Reuters. The data set consists of per barrel dollar spot prices from the 20th of May 1987 to the 23rd of February 2015, recorded at 16:30 UK time. In all hourly analysis we use a set containing 140 days of hourly Brent oil spot prices. The set is obtained from Sparebank1 Markets and contains hourly recordings from 02:00 the 28th of August to 09:00 the 12th of March.

We use the 3-month US Libor obtained from InterContinental exchange and 3-month NIBOR obtained from Norges Bank, both from 20th of May 1987 to 23rd of February 2015. We do not have information on the timing of these data sets.

## 3.2 Different types of crude oil

Three major types of crude oil dominate the market today: Brent, West Texas Intermediate (WTI) and Dubai/Oman. Brent is the referred to as North Sea oil, WTI is the main benchmark in the USA while Dubai/Oman is dominating in Asia. Brent and WTI oil are both light and sweet oils while the Dubai/Oman oil is heavier, sourer and is considered lower grade oil. In our study, we consider the Brent oil spot price. This for two reasons: Brent is the oil that is extracted from the North Sea and the Norwegian continental shelf and two-thirds of all oil contracts around the world is settled with Brent oil as the reference (Intercontinental Exchange 2013, U.S. Energy Information Association 2015). Figure 2 indicates that the prices of the two types of oil are highly correlated. We can also see that the spread between the two seem to have been increasing since around 2009. In addition, the trend up till 2009 was that the WTI was priced slightly above the Brent, while after 2009 it seems like Brent have been priced higher. The correlation between the two was 0,998 up till January 2009, while it was 0,78 after January 2009 and till today. The increase in spread may come from the startup of TransCanada Cushing Extension pipeline (U.S Energy Information Administration 2012). We omit any analysis of the Dubai/Oman because it is considered to be a different type of crude oil than the two other (Intercontinental Exchange 2013). In the following we always refer to the spot price of Brent crude oil when talking about the oil price.

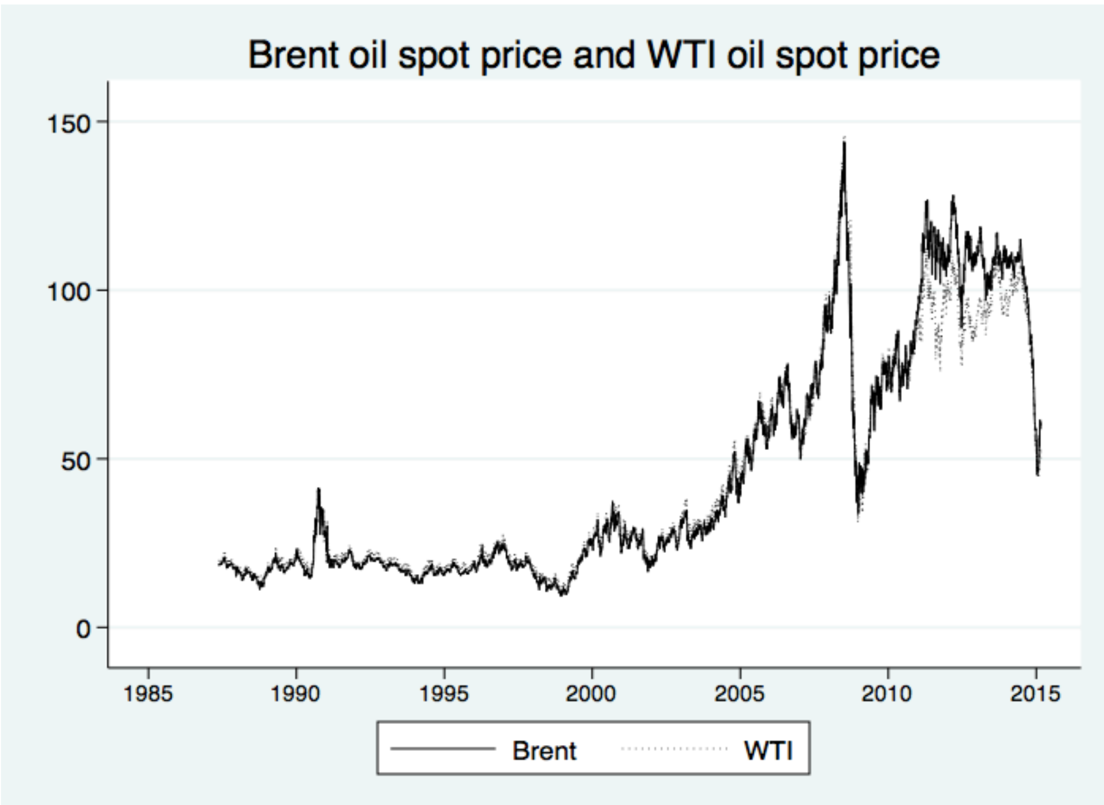


**Figure 2 – Time series of Brent oil spot price and WTI oil spot price**

Figure 2 suggests the presence of a shift around year 2000 from a time characterized by lower nominal prices and lower volatility to a time with higher and considerably more volatile nominal oil prices. A working paper by the European Central Bank (Fratzscher et al. 2014) recognizes that oil, since the early 2000s, increasingly has become a financial asset and attributes some of the rise in oil prices and oil price volatility to this phenomenon. The paper points to the development of a closer link between oil prices and other asset prices and oil prices are found to immediately reflect information of other asset prices. As examples, there is found a direct causal link between oil prices and exchange rates, and shocks to the return on equities are found to be important in explaining oil price movements. The same research finds that shocks to the financialization of the oil markets leads to a rise in oil prices. It is interesting that these findings were absent when analyzing data from before the 2000s, suggesting that the increased financialization of oil can account for a shift in behavior of oil prices (Fratzscher et al. 2014).

## 3.3 The exchange rate

Figure 3 illustrates a daily-recorded time series of the Norwegian exchange rate. It is illustrated in terms of Norwegian krone per unit of US dollar. A high value is therefore associated with a weak Norwegian krone. Increasing values is the same as depreciation of the Norwegian krone, while decreasing value is the same as appreciation of the Norwegian krone. As we can see, the exchange rate has sky rocketed (depreciated) from around 6,00 about a year ago, to 7,80 (while writing). This is a depreciation of 30% against the dollar. About halfway in this time series a change in the Norwegian Monetary policy was made. At the 29th of March 2001, a mandate was passed in the Norwegian Government giving Norges Bank the task to secure low and stable inflation at a long-term level of 2,5%. Before this, the monetary policy was supposed to secure a stable exchange rate relative to European currencies (Norges Bank 2015)[6]. When we refer to the exchange rate we mean the USD/NOK exchange rate unless otherwise is stated.

[6] http://www.norges-bank.no

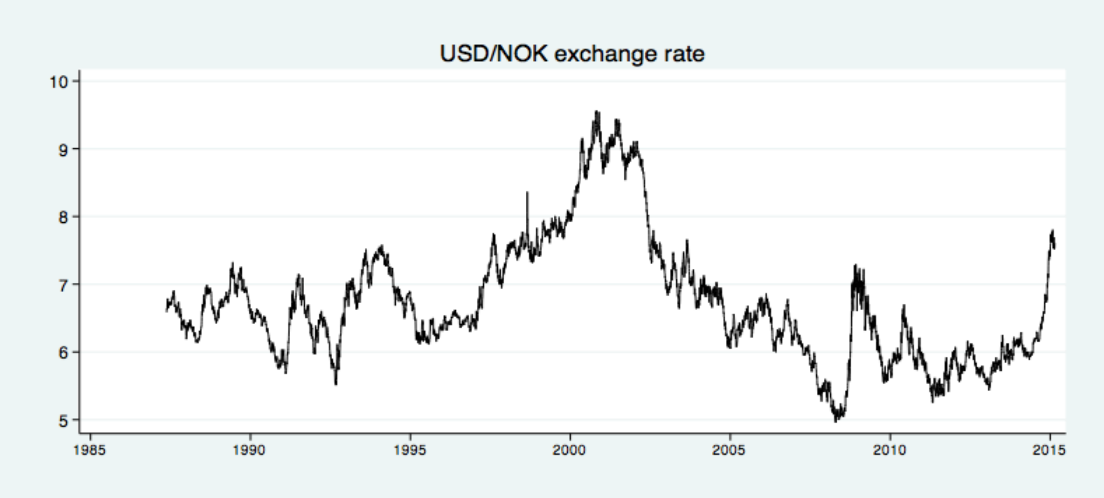


**Figure 3 – Time series of USD/NOK exchange rate**

## 3.4 First differenced logarithms

We calculate and use the natural logarithms of the exchange rates and oil prices in our analysis. This is useful as we, in our analysis, are interested in the growth rates of the series. By calculating the logarithms we reduce heteroscedasticity and mitigate extreme values. This helps satisfying the classical linear model assumptions. Using a non-stationary time series in a regression can lead to spurious results. In particular, the regression will not contain any long-run mean, as a result of permanent effects on the system from shocks to the variables. In addition, when non-stationary variables are used as input in a regression the usual test statistics will not follow their standard distributions (Brooks 2008).

Figure 4 and 5 plots the first differences of the time series for the log exchange rate and the log oil price. We can see that the graphs of the first differences do not display any significant trends and crosses their mean value frequently. This indicates that the first differences of both oil prices and exchange rates are stationary and that oil prices and exchange rates are integrated of order 1 (Brooks 2008).[7]

[7] When performing an augumented Dickey-Fuller (ADF) test in Stata 11 we clearly reject the null about the first differenced variable containing a unit root with t-values below -35 for both variables. For further explanation of the ADF test see Brooks (2008).

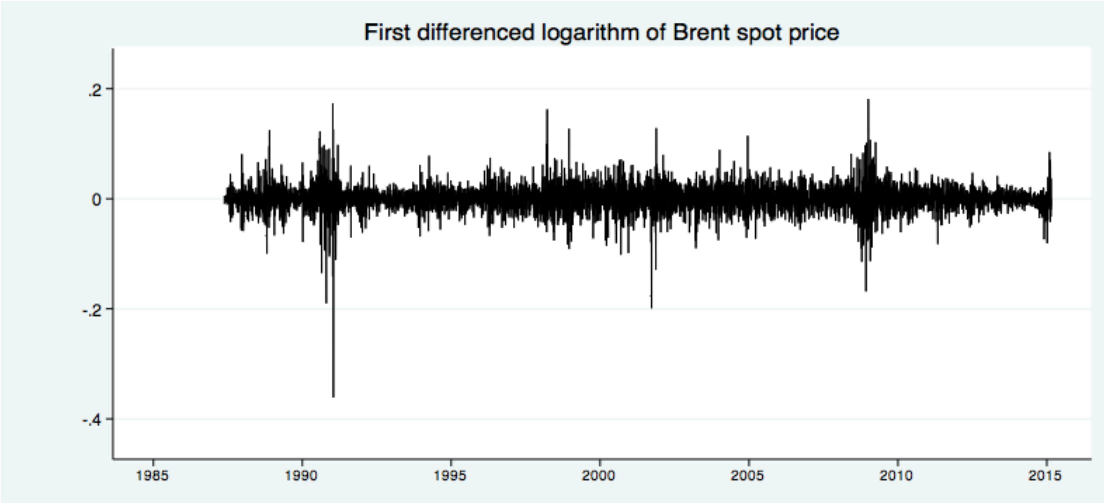


**Figure 4 – First differenced logarithm of Brent oil spot price**

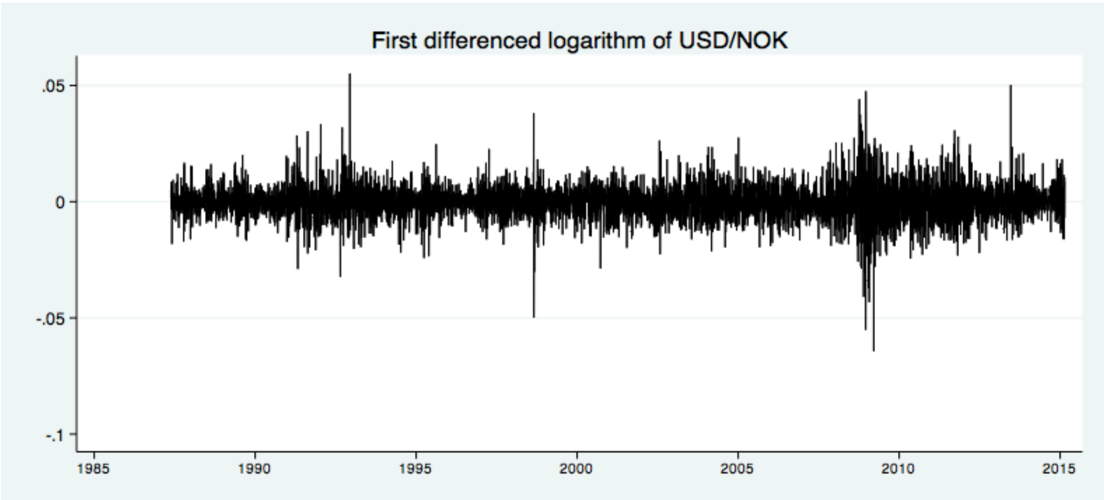


**Figure 5 – First differenced logarithm of USD/NOK exchange rate**

# 4 Regression results and the information overlap

Figure 1 shows the time series of both the Brent oil spot price and the Norwegian exchange rate. Note that, in this figure, we present the NOK/USD exchange rate, which denotes the dollar price per unit of the Norwegian krone, to better illustrate the relationship between the two variables. It is quite clear that the two parameters move closely, and that the relationship is much stronger in the last part of the set. In fact, if we look at the correlation between the two parameters (in level form) from the start of the data set to the $29^{th}$ of March 2001, the correlation is -0,16, which is weak and negative. If we look at the correlation between $29^{th}$ of March 2001 and till the end of the dataset, the correlation is 0,85, a very strong and positive relationship.

## 4.1 Empirical models and regression results

The analysis in this paper is based on simple models where the oil price is the only explanatory variable for the exchange rate. Reported in table 1 and 2 are the regression results from various empirical models when using both our daily time series of the exchange rate starting in 1987. Parameters are estimated using ordinary least squares (OLS) and the t-values for each parameters are reported in parentheses. Our two main empirical models are illustrated by equation 1 and 2. The former model coincides with the *contemporaneous model* while the latter coincides with the *true forecasting model.* $ex_t$ denotes the natural logarithm of the USD/NOK exchange rate and $p_t$ denotes the natural logarithm of the oil price. $\Delta$ denotes the first difference of the variables. Our discussion will focus on the results reported in table 1, where we use the data set from Bank of England. Results when using the data set obtained from Norges Bank, reported in table 2, yield the same conclusions.

$$\Delta ex_t = \beta_0 + \beta_1 \Delta p_t + u_t \tag{1}$$

$$\Delta ex_t = \beta_0 + \beta_2 \Delta p_{t-1} + u_t \tag{2}$$

From model 1 we estimate the value of $\beta_1$ to be -0,0519. This suggests that a 1 percentage point increase in the oil price growth rate leads to, on average, a $1 * 0{,}0519 = 0{,}0519$ percentage point decrease in the growth rate of the exchange rate. The constant is estimated to be 0,00003 and insignificant. This suggests no change in the exchange rate when the oil price change is zero. When looking at model 2 with lagged oil price changes, the relationship is much weaker. The estimated value of $\beta_2$ suggests that an increase of 1 percentage point in the lagged growth rate of the oil price, on average, leads to a decrease of $1 * 0{,}0076 = 0{,}0076$ percentage points in the growth rate of the exchange rate. Model 3 shows the results when including both contemporaneous and lagged oil price changes as explanatory variables.

$$\Delta ex_t = \beta_0 + \beta_1 \Delta p_t + \beta_2 \Delta p_{t-1} + u_t \tag{3}$$

The results show that the estimated $\beta$ coefficients do not change particularly, but the estimated $\beta_2$ becomes insignificant at the 5% significance level.

We further run regressions where the difference between the Norwegian and the US short-term interest rate are included as an explanatory variable. We found this relevant because of the well-known theory on uncovered interest parity (UIP). Briefly, it argues that the interest differential between two countries should equal the expected change in the exchange rate between the two countries (Chaboud & Wright 2003). The theory suggests that the nominal exchange rate will rise if the domestic interest rate rises and fall if the foreign interest rate rises. We therefore want to check if the interest rate differential between the US and Norway can offer any statistical explanatory power on the exchange rate between the two countries. $i_t$ denotes the interest rate spread between the 3-Month NIBOR and the 3-Month US LIBOR all dated at time $t$.[8] $\Delta$ denotes the first difference of the variables. Model 4 includes the contemporaneous interest rate spread as the only right-hand-side variable while model 5 only includes the lagged interest rate spread.

$$\Delta ex_t = \beta_0 + \gamma_1 \Delta i_t + u_t \tag{4}$$

$$\Delta ex_t = \beta_0 + \gamma_2 \Delta i_{t-1} + u_t \tag{5}$$

The estimated value of $\gamma_1$in model 4 is -0,0020. This indicates that if the contemporaneous spread increases by 1 unit (percentage point), the growth rate of the exchange rate will, on average, decrease by $0{,}002 * 100 = 0{,}2$ percentage points. This coincides with UIP. The estimated value of $\gamma_2$ suggests that if the lagged spread increases by 1 unit (percentage point), the growth rate of the exchange rate would, on average, decrease by $0{,}0003 * 100 = 0{,}03$

[8] The spread is defined as 3-Month NIBOR minus 3-Month US LIBOR, both variables in level form.

percentage points, but this coefficient is not significant. In table 1 and 2 we also report the estimated coefficients from the following empirical models:

$$\Delta ex_t = \beta_0 + \beta_1 \Delta p_t + \gamma_1 \Delta i_t + u_t \qquad (6)$$

$$\Delta ex_t = \beta_0 + \beta_2 \Delta p_{t-1} + \gamma_2 \Delta i_{t-1} + u_t \qquad (7)$$

Model 6 is a contemporaneous model with both change in log oil price and change in interest spread as right-hand-side variables. Model 7 includes lagged log oil price changes and lagged change in interest rate spread as right-hand-side variables. We observe that the inclusion of the spread barely changes the parameter of the oil price with regards to magnitude or t-value. This also holds for both model 6 and 7. Thus, controlling for the spread does not change our conclusions about the effect of the oil prices. Ferraro et al. (2015) considers daily interest rate differentials and concludes that they have no predictive power on the exchange rate. Since lagged values appear to be insignificant and because we are uncertain about the timing of our data on daily interest rates, we do not consider interest rate differentials any further. To summarize, negative contemporaneous oil price parameters suggest that an increase in the oil price is matched with a decrease in the exchange rate, as we would expect. Even more interestingly for our analysis is the significant negative values of the lagged oil price parameters, as this suggests an ability of previous oil price changes to explain future exchange rate changes.

**Table of regressions when using exchange rate data obtained from Bank of England**
- Oil price recorded at 16:30
- USD/NOK recorded at 16:00
- Dependent variable: $\Delta ex_t$

| | 1 | 2 | 3 | 4 | 5 | 6 | 7 |
|---|---|---|---|---|---|---|---|
| $cons$ | 0,00003<br>(0,35) | 0,00002<br>(0,25) | 0,00003<br>(0,35) | 0,00002<br>(0,22) | 0,00002<br>(0,23) | 0,00003<br>(0,33) | 0,00002<br>(0,25) |
| $\Delta p_t$ | -0,0519<br>(-14,40) | | -0,0517<br>(-14,34) | | | -0,0521<br>(-14,51) | |
| $\Delta p_{t-1}$ | | -0,0076<br>(-2,08) | -0,006<br>(-1,66) | | | | -0,0076<br>(-2,09) |
| $\Delta i_t$ | | | | -0,0020<br>(-6,18) | | -0,0021<br>(-6,41) | |
| $\Delta i_{t-1}$ | | | | | -0,0003<br>(-1,03) | | -0,0003<br>(-1,05) |

**Table 1 – Table of regressions: Exchange rate recorded at 16:00 (BOE)**

| Table of regressions when using exchange rate data obtained from Norges Bank.<br>- Oil price recorded at 16:30<br>- USD/NOK recorded at 13:30<br>- Dependent variable: $\Delta ex_t$ | | | | | | | |
|---|---|---|---|---|---|---|---|
| | 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| $cons$ | 0,00003<br>(0,32) | 0,00002<br>(0,29) | 0,00003<br>(0,35) | 0,00002<br>(0,24) | 0,00002<br>(0,24) | 0,00003<br>(0,3) | 0,00002<br>(0,29) |
| $\Delta p_t$ | -0,0341<br>(-9,28) | | -0,0334<br>(-9,11) | | | -0,0342<br>(-9,33) | |
| $\Delta p_{t-1}$ | | -0,0219<br>(-5,93) | -0,0208<br>(-5,68) | | | | -0,0219<br>(-5,93) |
| $\Delta i_t$ | | | | -0,0015<br>(-4,43) | | -0,0015<br>(-4,54) | |
| $\Delta i_{t-1}$ | | | | | -0,002<br>(-0,57) | | -0,0002<br>(-0,63) |

**Table 2 – Table of regressions: Exchange rate recorded at 13:30 (NB)**

## 4.2 Information overlap

One issue regarding the daily data sets is that the oil price is recorded *after* the exchange rate each day. This means that even when using lagged values, the change in oil price still overlaps the change in exchange rate. To illustrate:

- Using Bank of England's data: The change in oil price from 16:30 on Wednesday till 16:30 Thursday significantly explain some of the change in exchange rate from 16:00 on Thursday till 16:00 Friday – an overlap of approximately 30 minutes.
- Using Norges Bank's data: The change in oil price from 16:30 on Wednesday till 16:30 Thursday significantly explain some of the change in exchange rate from 13:30 on Thursday till 13:30 Friday – an overlap of approximately 3 hours.

Consequently, we cannot perform a regression with a true one-day lag, which in reality does not have an information overlap. We note that the regression results differ when changing the timing of the exchange rate data. This indicates that the timing is important. This importance is also evident when we later perform forecasting exercises; we therefore question why Ferraro et al. (2015) doesn't explicitly report the timing when analyzing daily data. It's also disappointing that we do not have access to a time series of the same length where there is no information overlap. However, we have access to 4-year data sets that enables us to illustrate the effect of timing. The results are illustrated in table 3. The reported t-statistics are from a regression of the lagged change in log oil price ($\Delta p_{t-1}$) on the change in log exchange rate ($\Delta ex_t$) recorded on different points in time (model 2). A t-statistic below -1,96 indicates a significant negative relationship between the two variables at a 5% significance level. We observe that as the information overlap decreases the relationship goes from being significant to being insignificant. We can also see that the coefficient decreases in value. This illustrates the importance of the timing of the data. This is also one of the reasons why we are not as optimistic as Ferraro et al. (2015) regarding our findings in chapter 6.

| **Regressions illustrating the importance of exchange rate data timing**<br>- Oil price recorded at 16:30<br>- Dependent variable: $\Delta ex_t$ recorded at different points in time<br>- Independent variable: $\Delta p_{t-1}$ | | | |
|---|---|---|---|
| Time | t-statistic | Coefficient | Overlap (hours) |
| 08:00 | -7,10 | -0,0967 | 8,5 |
| 09:00 | -6,16 | -0,0828 | 7,5 |
| 10:00 | -5,69 | -0,0766 | 6,5 |
| 11:00 | -4,91 | -0,0653 | 5,5 |
| 12:00 | -4,29 | -0,0586 | 4,5 |
| 13:00 | -3,88 | -0,0537 | 3,5 |
| 14:00 | -2,23 | -0,0305 | 2,5 |
| 15:00 | -0,28 | -0,0039 | 1,5 |
| 16:00 | -0,61 | -0,0083 | 0,5 |
| 17:00 | -0,69 | -0,0092 | 0 |
| 18:00 | -0,55 | -0,0074 | 0 |

**Table 3 – Regressions illustrating the importance of exchange rate data timing**

# 5 Exploiting the relationship - A trader's perspective

Motivated by the previous section, we construct a strategy where we go long the dollar if the oil price decreases and go short the dollar if the oil price increases. In spirit of the regression results in table 4 (section 5.1), the oil price change has to exceed a certain boundary before we execute a trade. For the first strategy we only take long positions if the oil price decreases more than a given boundary. We call this *Strategy 1*. Further, we construct a strategy where we trade on both increases and decreases in the oil price. With a decrease greater than a given boundary we go long the dollar and with an increase greater than a given boundary we go short the dollar. We call this *Strategy 2.* When trading we treat the currency as a stock with the price equal to the USD/NOK exchange rate. When going long the dollar we buy dollar bills in the market using NOK. When going short the dollar we “borrow” dollar bills and sell them in the market for NOK before buying them back with NOK at a later point in time. We always spend the whole portfolio of cash when trading.

The strategies involve boundaries to be crossed before we execute a trade. In strategy 1 we take a long position if the oil price decreases more than a given boundary. The boundary is formed as follows:

$$Long\ the\ dollar\ if\!: \frac{brent_t}{brent_{t-1}} < (1 - b_L) \tag{8}$$

$b_L$ denotes the boundary for going long. For strategy 2 we trade both on increases and decreases in the oil price. We impose two boundaries. One for going long (Equation 8) and one for going short:

$$Short\ the\ dollar\ if\!: \frac{brent_t}{brent_{t-1}} > (1 + b_S) \tag{9}$$

Using a rolling window of different sizes we see how our strategies are able to consistently outperform the buy and hold during different sub-sets of our sample. We also remove the last part of the set where the dollar appreciates almost 30% against the Norwegian krone in a couple of months to see if our strategies beat the buy and hold strategy under more “normal” conditions. Further we use hourly data to test our strategies with higher frequency trading. Our results show that when using a simple risk-adjusted measure similar to the Sharpe Ratio (1994) we are able to beat the buy and hold at both daily and hourly trading.

## 5.1 Implementation of boundaries

When using economic variables, changes in the independent variable might have to overcome a certain threshold before affecting the dependent variable; an idea presented by, among others, Dagenais (1969). When considering daily and hourly financial data, prices will fluctuate continuously even though there might not be any new relevant information in the market. Based on the same idea we want to increase the sophistication of our strategies by imposing boundaries for the changes in oil price to overcome before we execute a trade. Therefore we run a regression of the first differenced logarithm of the exchange rate on the first differenced logarithm of the oil price including a threshold variable. The regression results using the 4-year data set with zero information overlap yields interesting results.[9] Table 4 illustrates the results from the regression when including a threshold variable with an ex-post chosen optimal threshold.[10] The results illustrate that with no threshold variable, the basic relationship between oil price changes and exchange rate changes are negative, but not significant. When including a threshold variable, the threshold coefficient is significantly negative, while the coefficient on first differenced values is positive but not significant. This motivates us to trade when observing large changes in the oil price while doing nothing when observing small changes. The threshold model is illustrated by:

$$\Delta ex_t = \beta_0 + \beta_1 \Delta p_{t-1} + \beta_3 d_{1t-1} + u_t \tag{10}$$

The threshold variable, $d_{1t-1}$, can be interpreted as an observer of small shocks in the oil price:

$$d_{1t-1} = \begin{cases} \Delta p_{t-1} \; if \; \Delta p_{t-1} > \; 0{,}0213 \; or \; \Delta p_{t-1} < \; -0{,}0191 \\ 0 \; otherwise \end{cases} \tag{11}$$

This means that the log of the oil price must either increase by more than 2,13% or decrease by more than 1,91% for $d_{1t-1}$ to take on the value of $\Delta p_{t-1}$.

9 Exchange rate recorded at 17:00.

10 With optimal we mean the threshold that gives the most significant threshold variable (most negative t-value).

| Regression results: Threshold model<br>- Oil price recorded at 16:30<br>- USD/NOK recorded at 17:00<br>- 4-year data set with zero information overlap<br>- Dependent variable: $\Delta ex_t$ | | |
|---|---|---|
| | 10 | 2 |
| $cons$ | 0,0002<br>(0,78) | 0,0002<br>(0,97) |
| $\Delta p_{t-1}$ | 0,0419<br>(1,78) | -0,0082<br>(-0,61) |
| $d_{1t-1}$ | -0,0736<br>(-2,58) | |

**Table 4 – Regression results: Threshold model. Exchange rate recorded at 17:00**

## 5.2 Simulating the strategies on daily data

Because of the timing issue in our longest data sets we can't use this data when forming trading strategies that involve trading currency shortly after a signal given by the oil price. However, we can use the data set where the exchange rate is recorded daily at 17:00, half an hour after the oil price, to implement several trading strategies. This data set contains data from the 3rd of May 2010 to the 23rd of February 2015 with 1256 daily observations. Given our regression results our main strategy is simple. Oil price changes are defined as the change in price from 16:30 one day to the next. Exchange rate changes are defined as the change in the rate from 17:00 one day to the next. This means that we are able to trade the dollar once each day at 17:00. We construct two trading strategies and simulate their performance using the whole 4-year data set. During the period the dollar appreciated 28,16%. We compare our strategies with a buy and hold strategy of the dollar.

The data sets used ensure zero information overlap, which makes it valid for simulating the strategies performance. Using the whole window of observations to test our strategies is a somewhat dubious way of checking the performance. It is hard to beat the 30% appreciation of the dollar against the Norwegian krone during the last months. Therefore, we test our strategy on a rolling sample of observations. This means that we implement our strategies for a given window ($w$) of observations (days), roll this window over one step each time through the whole data set and report the results for each sub-sample.[11] The ending balance of cash from each trading period ($w$) is mathematically illustrated by the following equation:

[11] To illustrate: Say we choose a 50-day window. In this case we record the performance when implementing our strategies from day 1 to day 50, day 2 to day 51, day 3 to day 52 and so on through the whole sample.

$$C * \prod_{t=1}^{w} S_t \,, S_t = \begin{cases} \frac{E_t}{E_{t-1}} \, if \, \frac{brent_{t-1}}{brent_{t-2}} < (1 - b_L) \\ \frac{E_{t-1}}{E_t} \, if \, \frac{brent_{t-1}}{brent_{t-2}} > (1 + b_s) \\ 1 \; otherwise \end{cases} \tag{12}$$

$C$ denotes the starting amount of cash. $E_t$ and $brent_t$ denote the USD/NOK exchange rate and oil price at time $t$ respectively. By this we get a number of sub-sample returns equal to the total number of observations minus the window size. We do the same for the buy and hold strategy as a comparison. The buy and hold strategy is simply buying the dollar and holding it for the given window size. We perform the analysis on four different window ($w$) sizes: 50, 200, 500 and 700. The results are somewhat clear. The buy and hold strategy always have the highest average return, but also the highest standard deviation. If we use a risk-adjusted measure of performance, simply the average return divided by the standard deviation, strategy 1 and 2 often outperform the buy and hold strategy. Table 5 illustrates the results if we use a rolling window of 50 days. Strategy 1 and 2 are implemented with the ex-post chosen long ($b_L$) and short ($b_S$) boundaries of 1,8% and 1% respectively.[12] We keep these constant through the analysis. In table 5 we see that the buy and hold have the highest average return of 0,82%. Further we see that strategy 2 actually have the highest median. We can also observe that strategy 2 yields a positive return 806 out of 1207 times, while the buy and hold only yields positive return 643 times. The most important measure, where we correct for the risk taken, shows that both strategy 1 and 2 outperforms the buy and hold. We get more units of return per unit of risk taken. The reported t-values are the results from regression of the net change in cash (resulting from the trading strategies) on a constant. We do this to determine if the average return is significantly positive. A t-value above 1,96 indicates that a trading strategy have a significant positive expectation at the 5% significance level.[13]

[12] The boundaries are chosen based on giving the highest return for strategy 2 when implementing the strategies on the whole sample while also ensuring that frequent trading occur. With these boundaries the strategies give the most satisfying risk-adjusted returns while at the same time yielding high average returns.

[13] One can of course question whether the net results are independent observations (thus, question the use of the t-statistic) due to how they are constructed (through one-step rolling samples).

| | | | |
|---|---|---|---|
| Rolling window | 50 | | |
| Number of observations | 1207 | | |
| **Strategy** | Buy and hold | Strategy 1 | Strategy 2 |
| Average return | 0,82% | 0,37% | 0,63% |
| Median | 0,53% | 0% | 0,65% |
| Number of positive return | 643 | 557 | 806 |
| Standard deviation | 4,67% | 1,52% | 2,37% |
| Risk-adjusted measure | 0,18 | 0,24 | 0,27 |
| T-value | 6,11 | 8,52 | 9,20 |

**Table 5 – Strategy performance: Window size of 50 days**

In table 6 we see the results for a window size of 200 giving 1057 portfolio observations. The buy and hold yields an average return of 2,09% and a standard deviation of 8,82%. By this the risk-adjusted measure is 0,24, which is lower than for both strategy 1 and 2 with 0,35 and 0,44 respectively. We can also see that both strategies have a higher number of positive returns than the buy and hold.

| | | | |
|---|---|---|---|
| Rolling window | 200 | | |
| Number of observations | 1057 | | |
| **Strategy** | Buy and hold | Strategy 1 | Strategy 2 |
| Average return | 2,09% | 1,20% | 1,80% |
| Median | 2,61% | 1,48% | 1,96% |
| Number of positive return | 620 | 699 | 732 |
| Standard deviation | 8,82% | 3,44% | 4,10% |
| Risk-adjusted measure | 0,24 | 0,35 | 0,44 |
| T-value | 7,71 | 11,31 | 14,32 |

**Table 6 - Strategy performance: Window size of 200 days**

Increasing the window to 500 we see a difference in the results. With this window size we can see that the buy and hold outperforms our strategies with regards to both average return and risk-adjusted return. The buy and hold yields an average of 5,47% with a risk-adjusted return of 0,51. But, if we impose long and short boundaries of 2% and 1% respectively, strategy 1 actually outperforms the buy and hold strategy with a risk-adjusted return of 0,56.[14] Results are reported in table 7.

[14] Average return: 1,21%; Median: 0,72%; Number of positive returns: 512; Standard deviation: 2,17%; Risk-adjusted measure: 0,56.

| Rolling window | 500 | | |
|---|---|---|---|
| Number of observations | 757 | | |
| **Strategy** | Buy and hold | Strategy 1 | Strategy 2 |
| Average return | 5,47% | 1,14% | 1,91% |
| Median | 3,97% | 1,29% | 1,39% |
| Number of positive return | 497 | 467 | 515 |
| Standard deviation | 10,65% | 3,63% | 4,67% |
| Risk-adjusted measure | 0,51 | 0,31 | 0,41 |
| T-value | 14,12 | 8,65 | 11,24 |

**Table 7 – Strategy performance: Window size of 500 days**

When increasing the window to 700, strategy 2 is the best in terms of risk-adjusted performance with 0,77 units of return per unit of risk beating strategy 1 and the buy and hold with a risk-adjusted return of 0,41 and 0,62 respectively. Strategy 2 is also generates the highest number of positive returns. See table 8.

| Rolling window | 700 | | |
|---|---|---|---|
| Number of observations | 557 | | |
| **Strategy** | Buy and hold | Strategy 1 | Strategy 2 |
| Average return | 6,90% | 1,09% | 2,22% |
| Median | 7,32% | 1,41% | 1,75% |
| Number of positive return | 412 | 360 | 439 |
| Standard deviation | 11,06% | 2,69% | 2,89% |
| Risk-adjusted measure | 0,62 | 0,41 | 0,77 |
| T-value | 14,71 | 9,58 | 18,16 |

**Table 8 - Strategy performance: Window size of 700 days**

We further wanted to analyze how the results were affected if we stopped the dataset at 18th of June 2014 where the USD/NOK exchange rate was recorded at 6,0316.[15] After this date the dollar appreciates 26% against the Norwegian krone. This increase is considered to be an “outlier” in our data set. One might also discuss if this last increase is a state shift or structural break. If the exchange rate stabilizes around this level our strategies might be successful in the time coming. The results from the analysis are different. For window sizes 50 and 200, the buy and hold actually yields a negative expectation. Strategy 1 and 2 always yields positive expectation. At window sizes 500 and 700, the results are similar to earlier. At 500 the buy and hold outperforms all the strategies both in expected return and risk-adjusted return. Again, if we change the long and short boundaries to 2% and 1% respectively, strategy

[15] Delete the last 178 observations.

1 outperforms the buy and hold in terms of risk-adjusted returns.[16] The results are reported in table 9-12.

| Rolling window | 50 | | |
|---|---|---|---|
| Number of observations | 1029 | | |
| **Strategy** | Buy and hold | Strategy 1 | Strategy 2 |
| Average return | -0,26% | 0,34% | 0,64% |
| Median | -0,48% | -0,09% | 0,84% |
| Number of positive return | 465 | 456 | 692 |
| Standard deviation | 3,82% | 1,61% | 2,51% |
| Risk-adjusted measure | -0,07 | 0,21 | 0,25 |
| T-value | -2,15 | 6,79 | 8,16 |

**Table 9 - Strategy performance: Window size of 50 days (Without outlier)**

| Rolling window | 200 | | |
|---|---|---|---|
| Number of observations | 879 | | |
| **Strategy** | Buy and hold | Strategy 1 | Strategy 2 |
| Average return | 0,04% | 1,20% | 1,61% |
| Median | 0,06% | 1,48% | 1,41% |
| Number of positive return | 446 | 554 | 563 |
| Standard deviation | 6,87% | 3,74% | 4,42% |
| Risk-adjusted measure | 0,006 | 0,32 | 0,36 |
| T-value | 0,15 | 9,49 | 10,77 |

**Table 10 - Strategy performance: Window size of 200 days (Without outlier)**

| Rolling window | 500 | | |
|---|---|---|---|
| Number of observations | 579 | | |
| **Strategy** | Buy and hold | Strategy 1 | Strategy 2 |
| Average return | 1,37% | 0,78% | 0,60% |
| Median | 1,37% | -0,04% | 0,55% |
| Number of positive return | 320 | 289 | 337 |
| Standard deviation | 5,65% | 4,00% | 4,54% |
| Risk-adjusted measure | 0,24 | 0,20 | 0,13 |
| T-value | 5,84 | 4,70 | 3,20 |

**Table 11 - Strategy performance: Window size of 500 days (Without outlier)**

[16] Average return: 0,58%; Median: 0,72%; Number of positive returns: 334; Standard deviation: 2,00%; Risk-adjusted measure: 0,29.

| | | | |
|---|---|---|---|
| Rolling window | 700 | | |
| Number of observations | 379 | | |
| **Strategy** | Buy and hold | Strategy 1 | Strategy 2 |
| Average return | 2,31% | 2,37% | 2,42% |
| Median | 3,28% | 2,21% | 1,67% |
| Number of positive return | 412 | 360 | 439 |
| Standard deviation | 8,48% | 2,14% | 2,71% |
| Risk-adjusted measure | 0,27 | 1,11 | 0,89 |
| T-value | 5,29 | 21,56 | 17,37 |

**Table 12 - Strategy performance: Window size of 700 days (Without outlier)**

By using changes in the oil price as a predictor for changes in the exchange rate we are able to construct trading strategies that yield positive returns. In terms of risk-adjusted returns the strategies outperform a simple but intuitive benchmark: the buy and hold. This outperformance does not occur when looking at the 500-day window size, but we show that by changing the boundaries we are able to outperform the benchmark for this window as well. The conclusions hold both when including and excluding the latest sharp increase in the dollar that can be considered an outlier in our sample. Our results indicate that one can use the relationship between the variables to earn higher risk-adjusted profits than the buy and hold.

## 5.3 Simulating the strategies on hourly data

We also perform the same analysis on a 140-day long hourly data set from the 28th of August 2014 at 02:00 to the 12th of March 2015 at 09:00. This gives us 2987 hourly observations. The spot market for trading Brent is open from 02:00 to 23:00 on a trading day. The market for exchange rates is open all day long, but we have omitted the data between 23:00 and 02:00. Since the movements in the oil price mostly are very small from hour to hour, we tested a strategy similar to strategy 1, only with a boundary of zero. We call this *Strategy 3.* We perform exactly the same analysis as with daily data, but we do it on an hour-by-hour basis and with window sizes of 50, 250 and 600 hours.[17] Oil price and exchange rate changes are defined as hourly changes. We use preceding changes in the oil price to decide whether to take a position in the dollar. Since the changes are much smaller, we impose smaller boundaries. The boundaries for long and short are 0,031% and 0,047% respectively.[18] This analysis gives us a different result than with daily data. With the daily data, strategy 2

[17] It is critical to assume that we are able to buy the dollar exactly when the oil price confirms the signal, because the variables are recorded at the same point in time.

[18] The boundaries are ex-post chosen.

performed very well. This is not the case when using hourly data. Strategy 1 performs best. It must be mentioned that during the 140 days the dollar appreciated 32% against the Norwegian krone in a curve that is almost strictly increasing, as shown in figure 6. In a sample like this it is of course difficult to construct a strategy that increases more than the buy and hold when only trading in this particular asset.

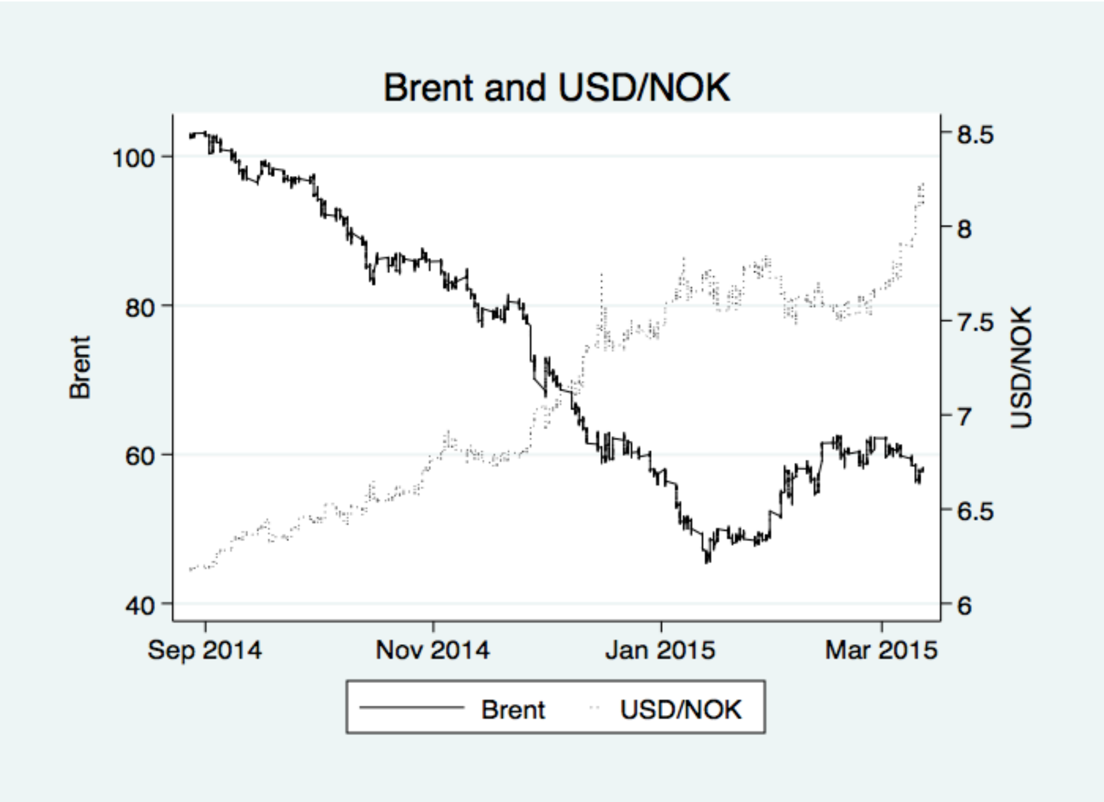


**Figure 6 – Time series of Brent oil spot price and USD/NOK exchange rate. 140-days hourly data**

Our results are very interesting and are reported in detail in table 13-15. Using a window of 50 hours, the buy and hold has the highest average return, but also the highest standard deviation. This makes the risk-adjusted performance worse than strategy 1 and 3, where strategy 1 is the best. The buy and hold yields positive profit in 2048 out of 2938 measures, closely followed by strategy 1 at 2028. With a window size of 250 hours strategy 1 and 3 outperforms the buy and hold at every measure but the average return. Regarding the risk-adjusted measure strategy 1 performs best while strategy 3 has most observations with positive return. If we further increase the window size to 600 hours strategy 1 performs very well. The average return is 4,88% while the buy and hold yields 4,94%, only 0,06 percentage points more, while the median of strategy 1 is higher. The risk-adjusted measure is also far better. If we look at the number of positive returns, the results are extraordinary. Strategy 1 gives positive return 2365 out of 2388 possible, which amounts to 99,04% of the observations.

| | | | | |
|---|---|---|---|---|
| Rolling window | 50 | | | |
| Number of observations | 2938 | | | |
| **Strategy** | Buy and hold | Strategy 1 | Strategy 2 | Strategy 3 |
| Average return | 0,474% | 0,390% | 0,299% | 0,377% |
| Median | 0,456% | 0,337% | 0,190% | 0,342% |
| Number of positive return | 2048 | 2028 | 1680 | 1987 |
| Standard deviation | 1,16% | 0,80% | 1,32% | 0,82% |
| Risk-adjusted measure | 0,41 | 0,49 | 0,23 | 0,46 |
| T-value | 22,16 | 26,46 | 12,77 | 24,90 |

**Table 13 - Strategy performance: Window size of 50 hours**

| | | | | |
|---|---|---|---|---|
| Rolling window | 250 | | | |
| Number of observations | 2738 | | | |
| **Strategy** | Buy and hold | Strategy 1 | Strategy 2 | Strategy 3 |
| Average return | 2,05% | 1,88% | 1,64% | 1,80% |
| Median | 1,87% | 1,56% | 1,22% | 1,40% |
| Number of positive return | 2245 | 2371 | 1757 | 2457 |
| Standard deviation | 2,32% | 1,87% | 3,86% | 1,90% |
| Risk-adjusted measure | 0,88 | 1,005 | 0,42 | 0,95 |
| T-value | 46,28 | 52,57 | 22,18 | 49,60 |

**Table 14 - Strategy performance: Window size of 250 hours**

| | | | | |
|---|---|---|---|---|
| Rolling window | 600 | | | |
| Number of observations | 2388 | | | |
| **Strategy** | Buy and hold | Strategy 1 | Strategy 2 | Strategy 3 |
| Average return | 4,94% | 4,88% | 4,36% | 4,61% |
| Median | 4,87% | 5,08% | 4,62% | 4,69% |
| Number of positive return | 2140 | 2365 | 1966 | 2334 |
| Standard deviation | 3,60% | 2,28% | 4,23% | 2,50% |
| Risk-adjusted measure | 1,37 | 2,14 | 1,03 | 1,84 |
| T-value | 67,02 | 104,24 | 50,38 | 67,02 |

**Table 15 - Strategy performance: Window size of 600 hours**

Through the whole data set of 2987 hours, the dollar increases 1560 times, only half of the set, while the oil price decreases 1525 times during the set. By this analysis, two things are somewhat clear: the buy and hold yields a higher expectation, but it also comes with higher risk. The conclusion for hourly trading is the same as for daily trading. We are able to construct strategies that outperform the buy and hold strategy in terms of risk-adjusted return no matter what window size we use.

## 5.4 Discussion on abnormal returns and the speed of adjustment

By blindly trading the dollar at signals given by the oil price we are actually able to make profits. The results suggest that our static trading strategy actually outperform a buy and hold strategy in terms of risk-adjusted performance, where our measure of risk is simply the standard deviation of the results. Our analysis is based on the assumptions of perfect capital markets where the market is frictionless (no trading costs or taxes), and where the participants in the market are price takers (Copeland et al. 2005). Since we have not imposed costs of trading such as a bid-ask spreads and transaction costs, this is just a theoretical, not an economical outperformance.

By our risk-adjusted measure it seems like we outperform the buy and hold strategy. But we are reluctant to say that our results "disprove" the weak form of efficiency suggested by the Efficient Market Hypothesis (EMH) (Fama 1970). The weak form of efficiency implicitly states that future prices of assets cannot be predicted by historical asset prices. This means that i.e. technical analysis should not work. Therefore, an investor should not be able to earn abnormal returns, typically measured by Jensens Alpha (Jensen 1967), by using a simple trading strategy such as ours. The question is therefore if our results can be considered as abnormal. To be abnormal they need to outperform a risk-adjusted benchmark. One example of a risk-adjusted benchmark for expected return for equity investors is the well-known Capital Asset Pricing Model by William F. Sharpe (1964) and John Lintner (1965) (as cited in Copeland et al. 2005). Abnormal returns can also be measured if we are able to construct a zero-beta portfolio, a portfolio with an expected return of zero. If this portfolio consistently yields positive returns, these can be considered as abnormal. However, when trading currency the risk-adjusted benchmark is not as clear. Melvin et al. (2013) stresses the issue of lacking a good benchmark for currency investing performance. For a fund manager investing in equities the benchmark might simply be the "market portfolio", typically S&P 500 or the Dow Jones Industrial Average. The same "market" does not exist for a currency trader, thus the possibility of a "passive" strategy is not apparent in the same way. Where the academics evaluate predictability measuring their forecasts against random walk using in example mean squared forecast error, traders measure their performance using a risk-adjusted measure such as the Treynor Index or the Sharpe Ratio (Melvin et al. 2013). We have used a measure similar to the Sharpe Ratio (1994) but without subtracting a benchmark rate of return (such as the risk free rate). This because it is not obvious what the proper risk free rate is, and it doesn´t change the ranking of our strategies.

One more thing should be noted regarding our trading strategies. Information spreads through the Internet extremely quickly and traders can get their orders executed in milliseconds when there is sufficient liquidity in the asset. This means that performing daily and hourly trading is relatively slow. In a study by Patell and Wolfson (1984) they used a simple trading strategy where they bought stocks where the earnings or dividend announcement exceeded what was expected by the "Value Line Investment Survey", and sold the stock short if the opposite happened. They concluded that there was activity in the stock price in the hours preceding the announcement, but that a very big amount of the price reaction came within five to ten minutes after the news. In the same paper they argue that the poorer the trading rule is, the faster the market will take advantage of any opportunities (Patell & Wolfson 1984). This indicates that less sophisticated strategies, such as ours, simulated on relatively low frequency data should have a hard time generating any profit. In a research by Dann et al. (1977) it is found that in a block trade (trade of 10 000 shares or more) the price drops significantly, but readjusts within 15 minutes of trading. These two studies coincide with the weak form of efficiency defined in EMH (Fama 1970). We emphasize that the studies are based on stock markets, but both oil and currency are financial assets that are traded extensively in the market. Today, information travels faster and thereby the market probably reacts more rapid. It is natural to assume that the same speed of adjustment to a large extent applies for currency markets as well. This makes us question if our "outperformance" occurs simply by choosing an irrelevant benchmark or by chance in our specific sample.
We believe that at strong an robust relationship is present, and due to the speed of adjustment we think that even higher profits could be generated by the implementation of our strategies on even higher frequency data such as minute- or second-data. Further, costs of trading have to be considered to make a judgment on the real profitability of our strategies. We also recognize that the boundaries are ex-post chosen and therefore will not necessarily yield satisfactory results in the time coming. The success of utilizing strategies such as ours relies heavily on the boundaries, implying that choosing correct ones *ex-ante* are crucial.

# 6 Forecasting ability – A statistician's perspective

Ferraro et al. (2015) mainly focuses on the ability of WTI oil price changes to predict changes in the Canadian dollar, but they also include a small sub-section on trying to predict the USD/NOK exchange rate using the WTI oil price. In our analysis we use more up-to-date data sets and several data sets for the exchange rate that differ with regards to timing, frequency and length. In addition, we consider the Brent oil price instead of the WTI oil price. We first consider daily data and perform an out-of-sample fit exercise using contemporaneous variables before conducting a true forecasting exercise. The true forecasting exercise involves the use of lagged oil price changes in an effort to explain exchange rate changes. We are able to find significant predictive ability and robust results for the contemporaneous model when it is compared against a random walk model. In addition, we find significant results for the true forecasting model if the rolling sample window is bigger than ¼ of the total sample size when using the data set from Norges Bank. This result does not hold when using the data set from Bank of England, where the information overlap is shorter. However, when letting the forecasting performance of our true forecasting model vary over time we find periods in the past where the model performed significantly better than the random walk. We point to the activity of traders in the financial markets with access to higher than daily frequency data and perform the same analysis using hourly data. We find significant results for the contemporaneous model but not the true forecasting model.

The forecasting model we evaluate is a very simple model. The change in oil price is the only included predictor for the change in exchange rate. This is the same model as considered in the analysis of Ferraro et al. (2015). In a survey of exchange rate models Rossi (2013) states that predictability is mostly apparent when one or more of the following hold: When the predictors are Taylor rule or net foreign asset fundamentals, when the model is linear and when a small number of parameters are estimated. Also, in favor of the simple model, Amano and Norden (1998) emphasizes that the oil price can be considered as an exogenous variable in the macroeconomic sense. They consider the link between the US real exchange rate and the real oil price. By pointing to the time series of the oil price they explain that the series is characterized by major supply-side shocks attributable to political conflicts in the Middle East, for which history offers no alternative explanation. This view, but for other countries real exchange rates, is supported by Chen and Rogoff (2002) as they state that primary commodities generally are exogenous to some small, but major commodity exporting countries. An explanation they suggest is that commodity products are transacted in highly centralized global markets. Obstfeld and Rogoff (1996) state that Canada is an open economy that is small with regards to the global oil market (as cited in Ferraro et al. 2015).[19] Accordingly, this makes it possible to assume that Canada is a price taker in that market and that the oil price might be considered an exogenous observable terms of trade shock. Since the description of the Canadian economy also characterize the Norwegian economy we assume that the argument regarding exogenity applies for Norway as well.

## 6.1 Two different exercises

When trying to create an economically valuable forecast, it indeed seems intuitive to use past changes in the oil price to predict future changes in the exchange rate. We call a model of this type a true forecasting model. If such a model fails, it might not be correct to dismiss the relationship between changes in the oil price and changes in the exchange rate of Norway. The failure of such a prediction model relies heavily on the relationship between past and future changes in the oil price, not only on the relationship between the oil price and the exchange rate. So, the failure might occur because yesterday's change in the oil price might not be a good predictor for todays change in the oil price (Ferraro et al. 2015). Further, it might be the case that yesterday's news is in fact yesterday's news, and that the market already has responded to yesterday's change because of the high trading activity in the two financial assets. This point highlights a possible need for even more frequent data than we consider here.

To more closely study the relationship between oil price changes and changes in the exchange rate we first conduct a contemporaneous forecasting analysis. This is a model where a prediction of the change in the exchange rate at time $t + 1$ is made by using the value of the change in the oil price at time $t + 1$. Actually, this is an out-of-sample fit exercise (Ferraro et al. 2015). To use such a model in practice one would have to wait until the date of the forecast to record the oil price, and then use this to "predict" an already realized change in the exchange rate. This might seem counter intuitive. However, if a model of this kind performs well (produces good out-of-sample fits) it may still be valuable in practical manners. Say that one has good predictions for tomorrows oil price, we could use the contemporaneous model to predict the change in tomorrow's exchange rate (Ferraro et al. 2015).

[19] We do not have access to the original work of Obstfeld and Rogoff (1996).

When presenting the results from the exercises we always report the DM-statistic for a range of different window sizes. Rossi and Inoue (2012) point to a common tendency in research to present empirical result for one window size and highlight two concerns regarding this. First, if the researcher tries only one window size, predictive ability might not be detected while it in reality exists for another window size. Second, the reported window size yielding satisfactory results may be a lucky shot found by chance after testing arbitrary window sizes. Changing the window size will affect how the sample is split into in-sample and out-of-sample parts and can lead to different empirical results (Rossi & Inoue 2012).

## 6.2 The Diebold and Mariano test

The test statistic we use to compare forecasts is the Diebold and Mariano statistic (DM statistic) (Diebold & Mariano 1995).[20] This statistic can be used as a formal statistical measure of the relative forecasting ability between two models and can determine whether one model generate significantly better forecasts than another.

Let us consider two different forecast series that can be compared to the true realized values. By subtracting the forecasted values from the actual realized values we obtain estimates of the forecast errors of the two forecast series. Let ${e_{1t}}_{t=1}^{T}$ and ${e_{2t}}_{t=1}^{T}$ be the forecast errors at time $t$, associated with the time series $y_t$, for model 1 and 2 respectively. The DM test aims to assess the loss associated with each forecast and makes assumptions directly on the forecast errors (Diebold and Mariano, 1995). The time $t$ loss associated with a forecast (say 1) can be seen as a function ($g$) of the true realized value ($y_t$) and the forecasted value ($\hat{y}_t$), $g(y_t, \hat{y}_{1t})$. Let us denote the loss differential series associated with each forecast as $d_t = [g(y_t, \hat{y}_{1t}) - g(y_t, \hat{y}_{2t})]$. According to Diebold and Mariano (1995) the critical assumption to make is that the loss differential is covariance stationary (DM assumption). This assumption can be summed up by the following assumptions:

$$E(d_t) = \mu, \forall\, t \tag{13}$$

[20] The main theory on the DM-statistic, as well as the mathematical derivations, are entirely based on the papers *Comparing Predictive Accuracy* by Diebold and Mariano from 1995 and *Comparing Predictive Accuracy, Twenty Years Later: A Personal Perspective on the Use and Abuse of the Diebold-Mariano Tests* by Diebold from 2013.

$$cov(d_t, d_{t-\tau}) = \gamma(\tau), \forall\ t \tag{14}$$

$$0 < var(d_t) = \sigma^2 < \infty \tag{15}$$

The null hypothesis we wish to test is one of equal forecast accuracy for forecast 1 and 2, which is equal to testing whether the population mean of the loss differential series is 0 (Diebold & Mariano 1995). This corresponds to the following null hypothesis, $H_0: E[d] = 0$. According to Diebold and Mariano (1995) this is an asymptotic test and standard results can be used to deduce the asymptotic distribution of the sample mean loss differential. The distribution can be illustrated by:

$$\sqrt{T}\left(\bar{d} - \mu\right) \xrightarrow{D} N\left(0, 2\pi f_d(0)\right), \tag{16}$$

where

$$\bar{d} = \frac{1}{T}\sum_{t=1}^{T}[g(e_{1t}) - g(e_{2t})], \tag{17}$$

is the sample mean loss differential,

$$f_d(0) = \frac{1}{2\pi}\sum\nolimits_{\tau=-\infty}^{\infty} \gamma_d(\tau), \tag{18}$$

is the spectral density of the loss differential at frequency 0 and

$$\gamma_d(\tau) = E[(d_t - \mu)(d_{t-\tau} - \mu)] \tag{19}$$

is the autocovariance of the loss differential at time difference $\tau$. $\mu$ is the mean loss differential of the population. We can see from the formula for $f_d(0)$ that the DM test recognizes that the forecast errors and hence the loss differential series may be serially correlated. This calls for a robust calculation of the standard error of the loss differential (Diebold 2013). More specifically, Diebold and Mariano (1995) argues that the sample mean loss differential can be considered as approximately normally distributed in large samples with mean $\mu$ and variance $\frac{2\pi f_d(0)}{T}$ and suggest the following large-sample $N(0, 1)$ statistic (DM) for testing the null hypothesis of equal forecast accuracy:

$$DM = \frac{\bar{d}}{\sqrt{\frac{2\pi \widehat{f_d}(0)}{T}}}, \tag{20}$$

where $\widehat{f_d}(0)$ is a consistent estimate of $f_d(0)$. Further, a consistent estimate of $2\pi f_d(0)$ is obtained by taking a weighted sum of the available sample autocovariances (Diebold & Mariano 1995):

$$2\pi \widehat{f_d}(0) = \sum_{\tau=-(T-1)}^{(T-1)} 1(\frac{\tau}{S(T)})\,\widehat{y_d}(\tau), \tag{21}$$

where

$$\widehat{y_d}(\tau) = \frac{1}{T}\sum_{t=|\tau|+1}^{T} (d_t - \bar{d})(d_{t-|\tau|} - \bar{d}). \tag{22}$$

Here, $1(\frac{\tau}{S(T)})$ denotes the lag window and $S(T)$ denotes the truncation lag.

To conclude, Diebold (2013) states that when the DM assumption holds, we have the following under the key null hypothesis:

$$DM = \frac{\bar{d}}{\widehat{\sigma_{\bar{d}}}} \xrightarrow{d} N(0,1), \tag{23}$$

$\widehat{\sigma_{\bar{d}}}$ denotes the consistent estimate of the standard deviation of $\bar{d}$ detailed in the discussion above. This leads us to the use of $N(0,1)$ critical values when using the DM-statistic for model comparison. This means that we reject the null of equal predictive accuracy at the 5% level if $|DM| > 1{,}96$.

Regarding the choice of lag window and truncation lag, Diebold and Mariano (1995) points to a familiar result that optimal k-step-ahead forecast errors are at most $k - 1$

dependent. Although they mention that $k-1$ dependence may be violated for many reasons, they argue that it seems reasonable to take $k-1$ dependence as a benchmark for a $k$-step-ahead forecast error. They suggest the use of the uniform, or rectangular, lag window, illustrated by:

$$1\left(\frac{\tau}{S(T)}\right)=\begin{cases}1 \ for \left|\frac{\tau}{S(T)}\right| \leq 0\\ 0 \ otherwise\end{cases} \tag{24}$$

In the practical application this means that $k-1$ sample autocovariances will be used in the estimation of $f_d(0)$ and $S(T)=k-1$. Because a uniform window assigns unit weight to all the covariances included the estimator will be consistent under $k-1$ dependence (Diebold & Mariano 1995).

### 6.2.1 Forecast evaluation with the Diebold and Mariano test

When testing both models we use the Diebold and Mariano statistic (DM-statistic) to test forecasting ability. This allows us to compare two forecasting models and evaluate if one performs statistically better than the other. As the comparable prediction model we always use a standard random walk "no-change" model. This means that the random walk forecast for changes in the exchange rate is zero, implying that the expectation of today's exchange rate is simply yesterday's exchange rate. The random walk model is recognized to be the toughest benchmark to beat by Ferraro et al. (2015) and Rossi (2013). For the measure of forecast errors we use mean squared errors (MSE). In all the cases below a test statistic less than -1,96 indicates that our model performs significantly better than a random walk model at a 5% significance level.

Intentionally, the DM-statistic was designed to evaluate forecasts that aren´t based on econometric models, so called model-free forecasts (Diebold 2013). However, Diebold (2013) recognizes that it has been common to use the DM-statistic to compare the forecasting ability of econometric models. Different approaches exist and to implement them the simple assumptions on the error loss differential are replaced by assumptions on the econometric models. Diebold (2013) emphasizes that that the approaches used may violate the DM assumption and that a number of aspects regarding the models should be considered. For example, when the models evaluated are nested the DM assumption enabling the researcher to use the DM-statistic with asymptotic $N(0,1)$ critical values, does not hold. The models we evaluate in this analysis are nested, something that is also recognized by Ferraro et al. (2015). At the same time he arguments that the DM-statistic can be used and points to the work of, among others, Giacomini and White (2006).[21] According to Diebold (2013), violations of the DM assumptions are often small, in which case the loss differential would be approximately stationary, and it seems to be a common view that the DM assumption may often be solicited without causing problems when comparing models. This speaks in favor of the procedure Ferraro et al. (2015) uses, which involves implementing the DM-statistic and using the asymptotic $N(0,1)$ critical values.

## 6.3 The contemporaneous model on daily data

First we analyze the forecasting performance of a contemporaneous model. The forecasts are all one-step-ahead based on daily observations. We use the following model from Ferraro et al. (2015):

$$\Delta ex_t = \beta_0 + \beta_1 \Delta p_t + u_t, t = 1, 2, \ldots, T \tag{25}$$

$ex_t$ and $p_t$ is the natural logarithm of the USD/NOK exchange rate and Brent oil spot price respectively, while the $\Delta$ denotes the first difference of the variables. $u_t$ is an error term that Ferraro et al. (2015) speaks of as "unforecastable". *t* represents point in time, $T$ represents the last observation. We use changes in the oil price to predict changes in the exchange rate, while both changes are recorded at the same day. Thus, this is more of an out-of-sample fit exercise rather than a true forecasting exercise. The forecasts are calculated as follows:

$$\Delta ex_{t+1}^f = \widehat{\beta_{0t}} + \widehat{\beta_{1t}} \Delta p_{t+1}, t = R, R+1, R+2, \ldots, T-1 \tag{26}$$

The $\hat{\beta}$ coefficients in (26) are estimated from a rolling sample of observations using regression (25) with a sample window of observations $\{t-R+1, t-R+2, \ldots, t\}$ where $R$ is the window size. $\Delta ex_{t+1}^f$ denotes our forecast (out-of-sample fit) for tomorrow´s change in

[21] Giacomini and White (2006) consider an unconditional test that is "(…) the Diebold and Mariano (1995) test extended to an environment that permits parameter estimation" (Giacomini & White 2006, p. 1571) and argue that the test can be used on both nested and non-nested models. An unconditional test compares forecasting models and evaluates which produces the most accurate forecasts on average.

the exchange rate. Using data on the exchange rate obtained from Bank of England the exercise gives us highly robust and significant results no matter what window size is used for the parameter estimation. This means that our prediction model outperforms the random walk model. The evidence of statistical significance is reported in table 16 where we can see that the DM-statistic is less than the 5% critical value using a wide range of window sizes, indicating that the null hypothesis of equal predictive ability is rejected. These results also hold when using data on the exchange rate obtained from Norges Bank. The DM-statistic for different window sizes is recorded in table 17. The results are somewhat stronger in the case with exchange rate data from Bank of England where the oil price and the exchange rate are recorded more closely in time.

**DM-statistics for different window sizes**
- Contemporaneous model compared against a random walk
- Exchange rate data recorded at 16:00 from Bank of England

| In-sample window size | In-sample window | DM-statistic | Forecast window |
|---|---|---|---|
| 1/20 | 362 | -6.859 | 6882 |
| 1/18 | 402 | -6.922 | 6842 |
| 1/16 | 453 | -7.008 | 6791 |
| 1/14 | 517 | -7.095 | 6727 |
| 1/12 | 604 | -7.136 | 6640 |
| 1/10 | 724 | -7.299 | 6520 |
| 1/8 | 906 | -7.487 | 6339 |
| 1/6 | 1207 | -8.079 | 6037 |
| 1/4 | 1811 | -8.563 | 5433 |
| 1/2 | 3622 | -10.17 | 3622 |
| 3/4 | 5433 | -10.74 | 1811 |
| 4/5 | 5795 | -9.761 | 1449 |

**Table 16 – DM-statistics. Contemporaneous model at different in-sample window sizes. Exchange rate recorded at 16:00 (BOE)**

| DM-statistics for different window sizes<br>- Contemporaneous model compared against a random walk<br>- Exchange rate data recorded at 13:00 from Norges Bank | | | |
|---|---|---|---|
| In-sample window size | In-sample window | DM-statistic | Forecast window |
| 1/20 | 362 | -4.743 | 6882 |
| 1/18 | 402 | -4.919 | 6842 |
| 1/16 | 453 | -5.160 | 6791 |
| 1/14 | 517 | -5.290 | 6727 |
| 1/12 | 604 | -5.394 | 6640 |
| 1/10 | 724 | -5.712 | 6520 |
| 1/8 | 906 | -5.998 | 6339 |
| 1/6 | 1207 | -6.633 | 6037 |
| 1/4 | 1811 | -7.252 | 5433 |
| 1/2 | 3622 | -8.519 | 3622 |
| 3/4 | 5433 | -9.073 | 1811 |
| 4/5 | 5795 | -7.585 | 1449 |

**Table 17 – DM-statistics. Contemporaneous model at different in-sample window sizes. Exchange rate recorded at 13:30 (NB)**

## 6.4 The true forecasting model on daily data

We now turn to the true forecasting model where lagged oil price changes are used to predict exchange rate changes. The forecasts we obtain and evaluate are all one-step-ahead forecasts based on daily observations. Following the same procedure as above, but now with lagged oil prices, we use a rolling window of different sizes to estimate the following model:

$$\Delta ex_t = \beta_0 + \beta_1 \Delta p_{t-1} + u_t, t = 1, 2, \dots, T \qquad (27)$$

In this case we get the following forecast model:

$$\Delta ex_{t+1}^f = \widehat{\beta_{0t}} + \widehat{\beta_{1t}} \Delta p_t, t = R, R+1, R+2, \dots, T-1 \qquad (28)$$

As before, $ex_t$ and $p_t$ denote the natural logarithm of the USD/NOK exchange rate and Brent oil price respectively. $\Delta$ denotes the first difference of the variables. $\Delta ex_{t+1}^f$ denotes our forecast for the $t+1$ change in the exchange rate, predicted by the change in the oil price at time $t$. $t$ represents point in time and $T$ represents the last observation. The coefficients in (28) are estimated from (27) with rolling sample window of observations $\{t-R+1, t-R+2, \dots, t\}$ where $R$ is the window size.

For forecasting purposes the model presented has a more intuitive look than the contemporaneous model. We tested our forecasting model using both data sets on the exchange rate dated back to 1987. When using the data set from Bank of England we are not able to find significant forecasting performance when attempting to beat a random walk. This is the case no matter what window size is used in the rolling sample regression. The DM-statistic for different window sizes is reported in table 18. One possible reason for this result is that daily data isn't frequent enough and doesn't contain valuable information because financial market participants utilize more frequent data. When using the data set from Norges Bank we are able to find significant predictive ability over the random walk for window sizes exceeding and including ¼ of the sample size. The DM-statistics for different window sizes are reported in table 19.

**DM-statistics for different window sizes**
- True forecasting model compared against a random walk
- Exchange rate data recorded at 16:00 from Bank of England

| In-sample window size | In-sample window | DM-statistic | Forecast window |
|---|---|---|---|
| 1/20 | 362 | 1.442 | 6882 |
| 1/18 | 402 | 1.099 | 6842 |
| 1/16 | 453 | 0.9236 | 6791 |
| 1/14 | 517 | 0.9625 | 6727 |
| 1/12 | 604 | 0.8074 | 6640 |
| 1/10 | 724 | 0.8356 | 6520 |
| 1/8 | 906 | 0.7560 | 6339 |
| 1/6 | 1207 | 0.8337 | 6037 |
| 1/4 | 1811 | 0.9146 | 5433 |
| 1/2 | 3622 | 0.2684 | 3622 |
| 3/4 | 5433 | 0.5186 | 1811 |
| 4/5 | 5795 | 1.329 | 1449 |

**Table 18 – DM-statistics. True forecasting model at different in-sample window sizes. Exchange rate recorded at 16:00 (BOE)**

| DM-statistics for different window sizes<br>- True forecasting model compared against a random walk<br>- Exchange rate data recorded at 13:30 from Norges Bank | | | |
|---|---|---|---|
| In-sample window size | In-sample window | DM-statistic | Forecast window |
| 1/20 | 362 | -0.3955 | 6882 |
| 1/18 | 402 | -0.4384 | 6842 |
| 1/16 | 453 | -0.4849 | 6791 |
| 1/14 | 517 | -0.6589 | 6727 |
| 1/12 | 604 | -0.8634 | 6640 |
| 1/10 | 724 | -1.135 | 6520 |
| 1/8 | 906 | -1.379 | 6339 |
| 1/6 | 1207 | -1.688 | 6037 |
| 1/4 | 1811 | -2.201 | 5433 |
| 1/2 | 3622 | -3.237 | 3622 |
| 3/4 | 5433 | -2.759 | 1811 |
| 4/5 | 5795 | -2.474 | 1449 |

**Table 19 - DM-statistics. True forecasting model at different in-sample window sizes. Exchange rate recorded at 13:30 (NB)**

One thing is important regarding these results. Since the exchange rate is recorded before the oil price each day, a model using a one time-unit lag will create an information overlap. Further, the differing results between the data sets indicate that the greater the overlap the more successful our model is, as discussed in section 4.2.

### 6.4.1 True forecasting performance of the threshold model

In chapter 5 we found that imposing boundaries, for which the oil price had to overcome before we interpreted it as a relevant signal, helped create a trading strategy that performed well. This was motivated by a simple regression of a threshold model (model 10). In table 20 below we can see the same regression with a different threshold variable (model 10a) when using the daily data set from Bank of England and the regression containing only the lagged oil price on the right side (model 2). The table shows that including the threshold makes the parameter on the lagged oil price insignificant and positive and that the parameter of the threshold variable is negative, although insignificant. The properties of the threshold variable, $d_{2t-1}$, is illustrated by the following:[22]

$$d_{2t-1} = \begin{cases} \Delta p_{t-1} \; if \; \Delta p_{t-1} > \; 0{,}0161 \; or \; \Delta p_{t-1} < \; -0{,}0151 \\ 0 \; otherwise \end{cases} \tag{29}$$

[22] The boundaries are ex-post chosen giving the lowest (most negative) t-value of the threshold coefficient.

We tested the forecasting performance of the threshold model on daily data. The results are reported in table 21. It is clear that the model does not significantly outperform the random walk model. We never reject the null about equal predictive ability. Ferraro et al. (2015) also tested the performance of a threshold model and state that it did not improve the forecasting performance. Due to the inability of the threshold model to improve forecasting ability in our case we do not consider it any further in the thesis.

| **Regression results: Threshold model**<br>- Oil price recorded at 16:30<br>- USD/NOK recorded at 16:00<br>- Dataset from 1987, Bank of England<br>- Dependent variable: $\Delta ex_t$ | | |
|---|---|---|
| | 10a | 2 |
| $cons$ | 0,00001<br>(0,15) | 0,00002<br>(0,25) |
| $\Delta p_{t-1}$ | 0,0136<br>(1,07) | -0,008<br>(-2,08) |
| $d_{2t-1}$ | -0,0231<br>(-1,74) | |

**Table 20 – Regression results: Threshold model. Exchange rate recorded at 16:00 (BOE)**

| **DM-statistics for different window sizes**<br>- Threshold forecasting model compared against a random walk<br>- Exchange rate data recorded at 16:00 from Bank of England | | | |
|---|---|---|---|
| In-sample window size | In-sample window | DM-statistic | Forecast window |
| 1/20 | 362 | 1,568 | 6882 |
| 1/18 | 402 | 1,251 | 6842 |
| 1/16 | 453 | 1,064 | 6791 |
| 1/14 | 517 | 0,9903 | 6727 |
| 1/12 | 604 | 0,6809 | 6640 |
| 1/10 | 724 | 0,7497 | 6520 |
| 1/8 | 906 | 0,6754 | 6339 |
| 1/6 | 1207 | 0,7617 | 6037 |
| 1/4 | 1811 | 0,8282 | 5433 |
| 1/2 | 3622 | -0,1767 | 3622 |
| 3/4 | 5433 | 0,3299 | 1811 |
| 4/5 | 5795 | 0,2195 | 1449 |

**Table 21 - DM-statistics. Threshold model at different in-sample window sizes. Exchange rate recorded at 16:00 (BOE)**

## 6.5 The forecasting ability changes over time

The analysis so far has used rolling regressions to estimate forecasts that have been evaluated based on the whole sample period. This means that all forecasts have been used to calculate one DM-statistic. Using daily data we got ambiguous results regarding the true forecasting model depending on the data set used. Rossi (2013) surveys exchange rate forecasting models over many years and reports that the forecasting performance of the models typically varies over time. Since we are considering a long time series over which the two variables have changed with regards to i.e. volatility, level and correlation (Fratzscher et al. 2014) it may be interesting to check whether the forecasting performance of our model has changed over the period. To do this we calculate the DM-statistic using a specified number of forecasts ($m$), less than the total number of available forecasts. We roll this calculation forward by one observation each step and calculate a total of $T - R - m$ unique DM-statistics. [23] By doing this we can see how the DM-statistic has changed over time. This procedure is the same as the one used by Ferraro et al. (2015).

For this test we use a rolling sample window ($R$) of 3622 observations and a window size ($m$) of 500 forecasts to calculate the DM-statistics. We perform the test on both of the data sets dated back to 1987 using the true forecasting model. The results using the data set from Bank of England are presented in figure 7. The figure indicates that our model performed significantly better than the random walk in a period between 2007 and 2009. During this time period the DM-statistic fluctuates around, mostly below, the critical value line. These results suggest that there was a time where the lagged forecasting model performed significantly better than a random walk model.

[23] To illustrate: Assume that we decide to check the performance of our model using the errors of 500 forecasts each time. The first DM-statistic is based on the first 500 of all available forecasts, the second DM-statistic is based 500 forecasts beginning with the second forecast, the third DM-statistic is based on 500 forecasts beginning with the third forecast and so on.

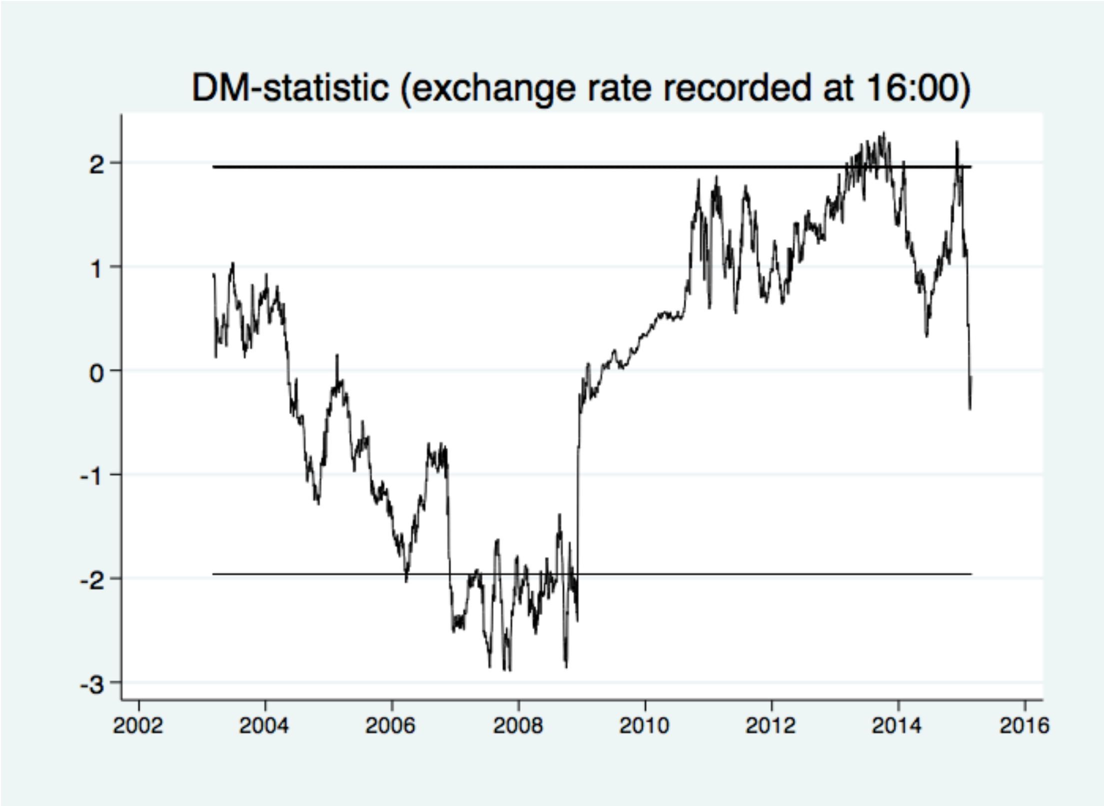


**Figure 7 – DM-statistic. True forecasting model. Performance over time. Exchange rate recorded at 16:00 (BOE)**

The results from the same test when using the data set from Norges Bank are reported in figure 8. Remember, using this data set we were able to find a significant DM-statistic when considering the total sample, as illustrated in table 19. Therefore, we would expect to observe more time periods indicating statistical outperformance of the random walk. In a period between 2005 and 2009 the DM-statistic is below the critical value line. In addition, there is a shorter period later in the sample where the same is observed. The figure indicates that the true forecasting model outperformed the random walk model during these periods. In general, for the two figures, we observe that the DM-statistic is mostly below zero indicating that the true forecasting model has lower forecasting errors than the random walk. Overall, the DM-statistics is typically at lower levels when using this data set from Norges Bank compared to when using the data set from Bank of England.

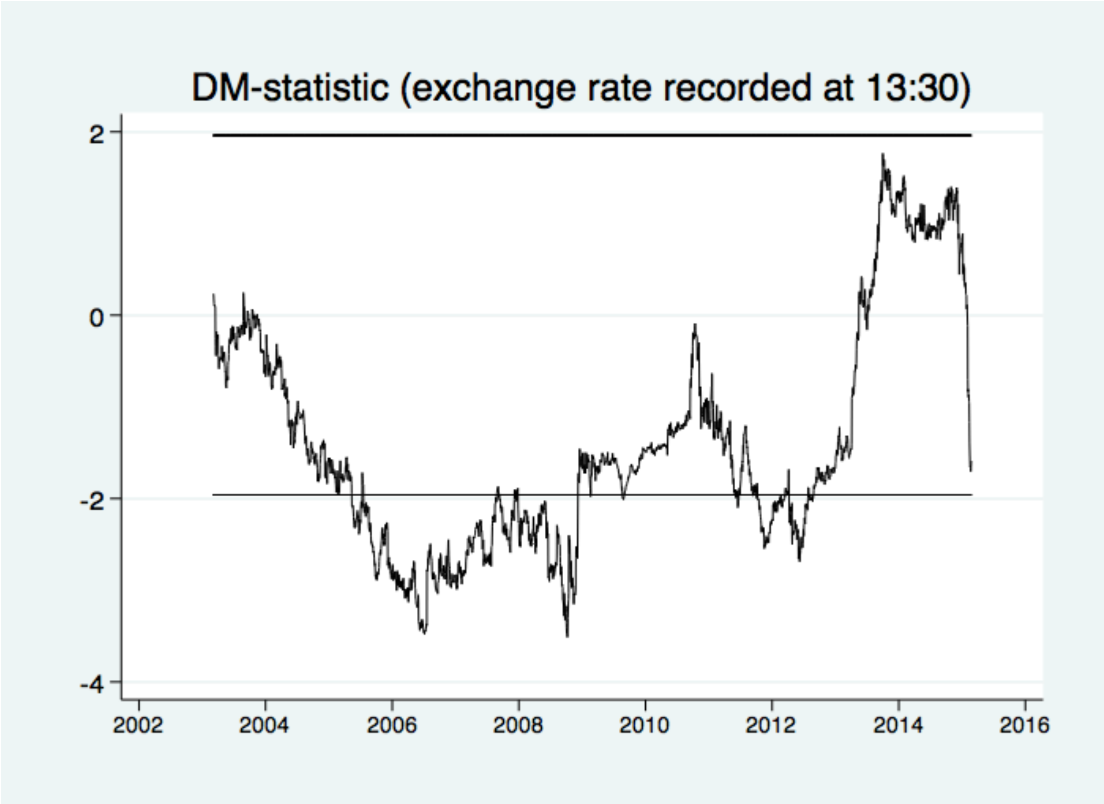


**Figure 8 - DM-statistic. True forecasting model. Performance over time. Exchange rate recorded at 13:30 (NB)**

## 6.6 Model performance on hourly data

It was pointed to earlier that daily prices might not be frequent enough to reflect information relevant for prediction of the exchange rate. One possible reason for this is that both oil and currency are financial assets that market participators, with access to high frequency data, trade continuously to exploit price discrepancies. This makes it interesting to test the models on hourly data using the 140-day hourly data set. Compared to the daily data sets this set covers a short time span, but still, the number of observations is 2987. We now consider hourly changes in the variables and all forecasts are one-step-ahead forecasts.

### 6.6.1 Contemporaneous model

To test the contemporaneous model we first do rolling regressions at a range of window sizes to estimate the following model:

$$\Delta ex_t = \beta_0 + \beta_1 \Delta p_t + u_t, t = 1, 2, \dots, T \tag{25}$$

Each forecast is calculated as follows:

$$\Delta ex_{t+1}^{f} = \widehat{\beta_{0t}} + \widehat{\beta_{1t}} \Delta p_{t+1}, t = R, R+1, R+2, \dots, T-1 \quad (26)$$

The models use the same notation as in the case with daily data, but Δ represents an hourly change, not a daily change. The results from the contemporaneous model are presented in table 22. We have identified significantly better forecasting performance for our model over a random walk. This result holds for all rolling sample window sizes up to and including half of the data set. The most significant result appears at a window size equal to ⅛ of the total sample with a DM-statistic equal to -3,34. The results are not as strong as when using daily data and do not hold for all in-sample window sizes.

| **DM-statistics for different window sizes**<br>- Contemporaneous model compared against a random walk<br>- 140 days with hourly data | | | |
|---|---|---|---|
| In-sample window size | In-sample window | DM-statistic | Forecast window |
| 1/20 | 149 | -2.531 | 2838 |
| 1/18 | 166 | -2.796 | 2821 |
| 1/16 | 187 | -2.819 | 2800 |
| 1/14 | 213 | -2.726 | 2774 |
| 1/12 | 249 | -2.790 | 2738 |
| 1/10 | 299 | -2.862 | 2688 |
| 1/8 | 373 | -3.053 | 2614 |
| 1/6 | 498 | -3.190 | 2489 |
| 1/4 | 747 | -3.335 | 2240 |
| 1/2 | 1494 | -2.966 | 1494 |
| 3/4 | 2240 | -1.663 | 747 |
| 4/5 | 2390 | -0.8808 | 597 |

**Table 22 – DM-statistics. Contemporaneous model at different in-sample window sizes. Hourly data**

### 6.6.2 True forecasting model

To test the forecasting model we first do rolling sample regressions at a range of window sizes to estimate the following model:

$$\Delta ex_t = \beta_0 + \beta_1 \Delta p_{t-1} + u_t, t = 1, 2, \dots, T \quad (27)$$

The forecasts are calculated as follows:

$$\Delta ex_{t+1}^{f} = \widehat{\beta_0} + \widehat{\beta_1}\Delta p_t, t = R, R+1, R+2, \dots, T-1 \tag{28}$$

The notation is the same as in the case with daily data, but again, the Δ represents hourly change. We discussed that data at increased time frequency might enable us to utilize valuable information not available in daily data, and therefore increasing the predictive ability of our model. Disappointingly, this is not what we find when using hourly data from our 140-day period. The results are presented in table 23. The DM-statistics are positive for all window sizes up to and including ⅙ of the sample set and negative for larger windows. None of the DM-statistics imply any significant difference in forecasting performance between our model and the random walk model. In other words we cannot reject the null hypothesis of equal predictive ability between the models. We also analyzed the forecasting performance when allowing for time-variation.[24] The results show that we are not able to reject the null stating equal predictive ability between our model and the random walk at any point in time.

**DM-statistics for different window sizes**
- True forecasting model compared against a random walk
- 140 days with hourly data

| In-sample window size | In-sample window | DM-statistic | Forecast window |
|---|---|---|---|
| 1/20 | 149 | 1.109 | 2838 |
| 1/18 | 166 | 0.9733 | 2821 |
| 1/16 | 187 | 0.9551 | 2800 |
| 1/14 | 213 | 0.8396 | 2774 |
| 1/12 | 249 | 0.6649 | 2738 |
| 1/10 | 299 | 0.2127 | 2688 |
| 1/8 | 373 | 0.1151 | 2614 |
| 1/6 | 498 | 0.0669 | 2489 |
| 1/4 | 747 | -0.2018 | 2240 |
| 1/2 | 1494 | -1.05 | 1494 |
| 3/4 | 2240 | -0.703 | 747 |
| 4/5 | 2390 | -1.147 | 597 |

**Table 23 – DM-statistics. True forecasting model at different in-sample window sizes. Hourly data**

With regards to our results a few things can be noted. First, market participants may trade away any forecasting opportunity in less than an hour by utilizing information available in even more frequent data, as already mentioned. In addition, we consider only one period of 140 days, where the oil price and the exchange rate both were characterized by rare behavior. The hourly data set spans from August 2014 to March 2015. What characterized the oil price

[24] Using $R$ = ½ of the total sample (1494) and $m$ = 500. We don´t report the results.

this period is a major fall from over $100 in August to under $60 in March. At the same time the USD/NOK exchange rate increased from around 6 NOK per dollar in August to over 8 NOK per dollar in March, indicating a major depreciation of the Norwegian krone. Although this is the kind of relationship between the two variables we have based our thesis on, it is worth noticing that the movements of the variables in the period that constitutes our sample is not characterized as normal.

## 6.7 The “dollar effect”

The dollar exchange rate and the Brent oil price are heavy drivers in the global economy and can be considered leading economic variables and important determinants of international trade (Ghalayaini 2011). Our analysis investigates the effect of Brent oil price changes on changes in the Norwegian krone exchange rate. The fact that the USD/NOK exchange rate is a dollar exchange rate, and the fact that the Brent oil price is denominated and settled in dollars makes the link between the oil price and the dollar exchange rate an important issue (Ghalayaini 2011). As the dollar changes relative to other currencies, the price of oil for holders of these currencies will change. This can influence the quantity of oil demanded on the global market place and lead to price changes in the dollar price of oil. Changes in demand for oil comes with changes in demand for US dollars, something that can lead to changes in the value of the dollar (Ghalayaini 2011). Remember, when the dollar changes on a general level the Norwegian krone relative to the dollar changes in value as well. This means that the highly correlated relationship between the USD/NOK variable and the oil price variable might be because of both variables being denominated in dollars. This issue is also recognized by Ferraro et al. (2015) and named the “dollar effect”. Statistical regressions show a strong significant relationship between changes in the oil price and changes in the USD/NOK exchange rate and this may be a result of the dollar effect, not a causal relationship between the Brent oil spot price and the Norwegian currency. To address this issue we use the same approach as Ferraro et al. (2015). We replace the USD/NOK exchange rate with the GBP/NOK exchange rate and conduct the contemporaneous and true forecasting exercise. When doing this we test if we can find the same results as found in our analysis above, when using the Norwegian currency relative to another currency than the dollar. For the GPB/NOK variable we only have one data set obtained from Bank of England with daily recordings of the exchange rate at 16:00 UK time. We expect that if the Norwegian krone is highly dependent on the Brent oil price, the NOK should depreciate against the GBP during an oil price fall and appreciate during an oil price increase. This unless there is substantial offsetting economical mechanisms between Norway and Great Britain during such times, but we do not have any reason to believe that this is the case. We use the following model:

$$\Delta exGBP_t = \beta_0 + \beta_1 \Delta p_t + u_t, t = 1, 2, \dots, T \tag{30}$$

$exGBP_t$ denotes the natural logarithm of the GBP/NOK exchange rate while the other notations are the same as before. The contemporaneous out-of-sample forecast model can then be illustrated by:

$$\Delta exGBP^f_{t+1} = \widehat{\beta_0} + \widehat{\beta_1} \Delta p_{t+1}, t = R, R+1, R+2, \dots, T-1 \tag{31}$$

The results from this out-of-sample fit exercise give the same results as the model with USD/NOK exchange rate; the DM-statistics are reported in table 24. The contemporaneous model forecasts better than the random walk model no matter what window size we use. This increases the validity of our findings of a robust and significant relationship between the Brent oil spot price and the Norwegian currency.

| **DM-statistics for different window sizes**<br>- Contemporaneous model compared against a random walk<br>- Exchange rate defined as GBP/NOK | | | |
|---|---|---|---|
| In-sample window size | In-sample window | DM-statistic | Forecast window |
| 1/20 | 362 | -3.118 | 6882 |
| 1/18 | 402 | -3.269 | 6842 |
| 1/16 | 453 | -3.350 | 6791 |
| 1/14 | 517 | -3.291 | 6727 |
| 1/12 | 604 | -3.221 | 6640 |
| 1/10 | 724 | -3.417 | 6520 |
| 1/8 | 906 | -3.837 | 6339 |
| 1/6 | 1207 | -4.409 | 6037 |
| 1/4 | 1811 | -5.017 | 5433 |
| 1/2 | 3622 | -5.789 | 3622 |
| 3/4 | 5433 | -5.977 | 1811 |
| 4/5 | 5795 | -6.08 | 1449 |

**Table 24 - DM-statistics. Contemporaneous model at different in-sample window sizes. GBP/NOK exchange rate recorded at 16:00 (BOE)**

When using lagged values of the oil price to forecast the exchange rate we do not find significant DM-statistics no matter what window size is used. The DM-statistics are reported in table 25 and indicates that the forecasting model cannot beat the forecasting performance of a random walk model. These results are also in line with what we have found when considering the USD/NOK exchange rate. We use the following model where the notations are the same as mentioned:

$$\Delta exGBP_t = \beta_0 + \beta_1 \Delta p_{t-1} + u_t, t = 1, 2, \ldots, T \quad (32)$$

Forecasts are calculated as follows:

$$\Delta exGBP^f_{t+1} = \widehat{\beta_0} + \widehat{\beta_1} \Delta p_t, t = R, R+1, R+2, \ldots, T-1 \quad (33)$$

| **DM-statistics for different window sizes**<br>- True forecasting model compared against a random walk<br>- Exchange rate defined as GBP/NOK | | | |
|---|---|---|---|
| In-sample window size | In-sample window | DM-statistic | Forecast window |
| 1/20 | 362 | 1.601 | 6882 |
| 1/18 | 402 | 1.428 | 6842 |
| 1/16 | 453 | 1.431 | 6791 |
| 1/14 | 517 | 1.342 | 6727 |
| 1/12 | 604 | 1.019 | 6640 |
| 1/10 | 724 | 1.155 | 6520 |
| 1/8 | 906 | 0.523 | 6339 |
| 1/6 | 1207 | 1.216 | 6037 |
| 1/4 | 1811 | 1.014 | 5433 |
| 1/2 | 3622 | 1.011 | 3622 |
| 3/4 | 5433 | 1.007 | 1811 |
| 4/5 | 5795 | 0.697 | 1449 |

**Table 25 – DM-statistics. True forecasting model at different in-sample window sizes. GBP/NOK exchange rate recorded at 16:00 (BOE)**

Ghalayini (2011) analyses the causal relationship between the price of oil and the dollar exchange rate and aims to investigate if there is interdependence between these variables. He concludes that even though oil prices are expressed in dollars, the changes in the dollar exchange rate have no significant effect on oil prices. Our findings regarding the “dollar effect” are in line with the results reported by Ferraro et al. (2015).

## 6.8 Chapter summary

We find that the contemporaneous model outperforms the random walk in a forecasting exercise based on the DM-test. This holds for a range of window sizes when utilizing the whole sample set. This is actually an out-of-sample fit exercise and these results cannot be used to make trading decisions in real life. We emphasize that if a good predictor of tomorrows oil price is found the contemporaneous model may be useful in practice (Ferraro et al. 2015). Also, the results indicate a strong and robust relationship between oil price changes and changes in the exchange rate.

Regarding the true forecasting model based on lagged changes of the oil price as the predictor we also find signs of predictive ability. However, the forecasting performance of our model depends on the timing of the data and changes a lot over our time series. Using the data set obtained from Bank of England we are not able to find significant results when using the whole set of available forecasts. When allowing for time variation in the performance of the model we identify periods where our model outperformed the random walk. When using the data set obtained from Norges Bank we find evidence of significant forecasting performance over the random walk when using the whole set of available forecasts. This holds for all rolling sample windows larger than ¼ of the total sample size. The different results between the two daily data sets may come from the fact that the lagged oil price changes overlaps the changes in the exchange rates by a varying extent.[25] We also conduct a simple analysis as an attempt to control for the dollar effect. The results do not change our main conclusions.

When increasing the data frequency to hourly observations we find evidence of significantly better performance over the random walk model when testing the contemporaneous model. The true forecasting model is not able to beat the random walk model at any point in time.

[25] Using the data set from Bank of England the overlap is half an hour and when using the data set from Norges Bank the overlap is three hours.

# 7 Two different perspectives – Comparing the results

We have analyzed to what extent the Brent oil price can be used to predict the USD/NOK exchange rate through both a trader's perspective and a statistician's perspective. The results from the two analyses are ambiguous. Our static trading strategies have mostly beaten the buy and hold strategy in terms of our risk-adjusted performance measure. This benchmark is an intuitive and natural first obstacle for any trader to overcome. Further, from a statistician's perspective, we take on the task of forecasting the exchange rate using the methodology of Ferraro et al. (2015). We find that when performing a true forecasting exercise using lagged values, we are only able to beat the random walk in a small sub-set of the data. Disappointingly, most of the time we must say that our model and the random walk perform statistically equal in terms of forecasting the exchange rate. Even when we increase the frequency to hourly data we cannot statistically outperform the random walk. We must mention that using the contemporaneous model performing the out-of-sample fit exercise, we find results that illustrate a strong and robust relationship between the two variables. Ferraro et al. (2015) argues the importance of these findings.

It is important to distinguish between the two perspectives. Traders are typically interested in risk-adjusted performance of a portfolio like the Sharpe Ratio while statisticians focus on evaluation of forecasting ability through the use of forecast error measures (Melvin et al. 2013). When choosing a model statisticians most often take on the task of beating the random walk, which research recognizes as the toughest benchmark to beat (Ferraro et al. 2015, Rossi 2013). Melvin et al. (2013) say, on the other hand, "*(...) beating a random walk is not a very useful evaluation metric for currency investing*" (Melvin et al. 2013, Abstract). Their paper stresses the importance of the difference between typical statistical measures, such as striving to generate lower forecast errors than the random walk, and the performance of an investment portfolio. When constructing an investment portfolio the object is to consistently generate positive returns, not necessarily to predict accurate level forecast of the exchange rate. Rather, correctly ordering the forecasts of future returns relative to one another is the critical task. Melvin et al. (2013) illustrates with a technical example how a simple trading decision based on forecasts can generate a profit through correctly *ranking* the returns of two exchange rates. Both return forecasts turned out to be substantially wrong in magnitude and one of the forecasts was wrong in terms of direction. In addition, the mean squared forecasting errors were larger than those of the random walk model.[26]

In our exercise we decide to go either long or short in the dollar based on signals from the oil price. We will earn money if we more or less consistently predict the right *direction* of the exchange rate change. It is irrelevant if we beat the random walk model or not, in terms of making money. As an analogy to the Melvin et al. (2013) example, for us correct ranking and decision-making implies correctly forecasting the direction of the exchange rate, not the magnitude of change. We will consistently earn profits if we are able to perform this task. In the statistical exercise of beating the random walk the magnitude of change is, on the other hand, critical.

We can draw different conclusions depending on the perspective one operates from. From the traders perspective we are effectively testing the true forecasting model and find that we are able to predict the direction of future exchange rate movements to some extent. By exploiting the link between the two variables we are able to form a trading strategy that outperforms a simple buy and hold strategy in terms of risk-adjusted returns, although we question our choice of the buy and hold strategy as benchmark.

From the statistical perspective our results are not as promising when it comes to the true forecasting model. Different data sets yield different conclusions and when testing the 4-year data set with no overlap the true forecasting model never outperforms the random walk.[27] However, we do find a strong and robust relationship between the oil price and the USD/NOK exchange rate documented by the results from the contemporaneous model. We suggest, in line with Ferraro et al. (2015), that if one had good predictions for future oil prices one could use the contemporaneous model to forecast exchange rate movements.

[26] For a more thorough walk-through of the example see ”*Forecasting Exchange Rates: An Investor Perspective”* (Melvin et al. 2013)

[27] We do not report the results from this DM test.

# 8 Conclusion

The recent behavior of the oil price and the exchange rate of Norway strongly motivated us to explore the relationship between the two variables. We have demonstrated what we think is an interesting relationship between them through the use of data series of different lengths at both daily and hourly frequencies. By focusing on the short-term predictive ability of the Brent oil price on the exchange rate of Norway, from the perspectives of both traders and statisticians, we have contributed to the vast literature on short-term exchange rate forecasting.

From the perspective of an investor wanting to earn a profit from trading the exchange rate we have constructed trading strategies and tested them on historical data. All the strategies constructed are based on the same idea: Long the dollar if the oil price decreases and short the dollar if the oil price increases. We find that implementing boundaries for the oil price changes to cross before executing a trade is a necessity for generating satisfactory profits when simulating on daily data, while not being of the same importance when simulating on hourly data. We conclude that we are able to outperform a simple buy and hold strategy in terms of risk-adjusted profits, but question the choice of the buy and hold strategy as a benchmark. In general, we are able to generate positive returns (and generate strategies with positive expectation), but these returns vary and are sometimes small in economic terms. Also, we have not imposed costs of trading which makes this a pure theoretical return. We will therefore not recommend any trader to perform the presented strategies on as low frequency data as ours. However, we suggest that our findings indicate the existence of a relationship between the variables that could be better exploited by more sophisticated trading models. Also, for research purposes we would like to simulate the strategies on longer time series and even higher frequency data.

From the perspective of a statistician we have evaluated the forecasting performance of simple models where oil price changes is the only explanatory variable for exchange rate changes. By adopting the methodology of Ferraro et al. (2015) two main models are evaluated against a random walk model in a forecasting exercise based on the Diebold and Mariano test (Diebold & Mariano 1995). The *contemporaneous model* significantly outperforms the random walk model. This holds no matter what in-sample window size is used for model estimation, for daily and hourly data and does not depend on timing of the exchange rate data.

The *true forecasting model* tests the ability of one-period lagged oil price movements to predict exchange rate movements. When evaluating the forecast performance of this model using daily data we find that the results are dependent on data timing. When the information overlap is three hours the model is found to significantly outperform the random walk model at large in-sample window sizes. With an information overlap of half an hour the model is not able to significantly outperform the random walk. However, we do identify periods in the past where it did outperform the benchmark. When evaluating hourly forecasts of the true forecasting model we found no significant outperformance of the random walk benchmark. Like Ferraro et al. (2015) we have recognized the dollar-effect issue and tested for it by conducting the same exercises while substituting the dollar exchange rate with the British pound exchange rate. It did not affect our main results.

The results from the true forecasting analysis show that the model performance is highly dependent on the timing of the data. Even at hourly data, with no information overlap, we document that there is no predictive ability. This indicates that we cannot use oil price movements alone to predict daily or hourly exchange rate changes one-step-ahead. We point to high frequency trading in the financial markets and desire even more frequent data to test this simple model on. Still, with an information overlap of three hours, we detect significant outperformance of the random walk. This and the results from our contemporaneous analysis document a strong and significant relationship between the variables, in line with the conclusions drawn by Ferraro et al. (2015). Further, we emphasize the possible practical usefulness of the contemporaneous model if one were to obtain good forecasts of the oil price.

## Reference list

Amano, R. A. and S. van Norden (1998): ”Exchange Rates and Oil Prices”, *Review of International Economics,* 6(4), 683-694.

Amano, R. A. and S. van Norden (1998): ”Oil Prices and the Rise and Fall of the US Real Exchange Rate”, *Journal of International Money and Finance,* 17, 299-316.

Bergo, J. (2004): ”Fleksibel inflasjonsstyring”, *Penger og Kreditt,* 2, 76-83.

Brooks, C. (2008): *Introductory Econometrics for Finance*. 2. edition. New York: Cambridge University Press.

Chaboud, A. P. and J. H. Wright (2003): ”Uncovered Interest Parity: It works, but not for long”, IFDP Working Papers No. 752, International Finance Discussion Papers.

Chen, Y. And K. S. Rogoff (2002): ”Commodity Currencies”, *Journal of International Economics,* 60(1), 133-160.

Chen, Y., K. S. Rogoff, B. Rossi (2010): ”Can Exchange Rates Forecast Commodity Prices?”, *The Quarterly Journal of Economics,* August, 1145-1194.

Cheung, Y., M. D. Chinn and A. G. Pascual (2005): ”Empirical exchange rate models of the nineties: Are any fit to survive?”, *Journal of International Money and Finance,* 24, 1150-1175.

Copeland, T. I., J. F. Weston and K. Shastri (2005): *Financial Theory and Corporate Policy.* 4. edition. United States of America: Pearson Education (US).

Dagenais, M. G. (1969): ”A Threshold Regression Model”, *Econometrica,* 37(2), 193-203.

Dann, L. Y., D. Mayers and R. J. Raab, Jr. (1977): ”Trading Rules, Large Blocks and the Speed of Price Adjustment”, *Journal of Financial Economics,* 4, 3-22.

Diebold, F. X. (2013): ”Comparing Predictive Accuracy, Twenty Years Later: A Personal Perspective on the Use and Abuse of Diebold-Mariano Tests”, Prepared for *Journal of Business and Economic Statistics* Invited Lecture, Allied Social Science Association Meetings, Philadelphia, January 2014.

Diebold, F. X., and R. S. Mariano (1995): ”Comparing Predictive Accuracy”, *Journal of Business and Economic Statistics,* 13, 253-265.

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

Ferraro, D., K. S. Rogoff and B. Rossi (2015): ”Can Oil Prices Forecast Exchange Rates?”, Economics Working Papers No. 1461, Department of Economics and Business, Universitat Pompeu Fabra.

Fratzscher, M. D. Schneider and I. Van Robays (2014): ”Oil prices, exchange rates and asset prices”, Working Paper Series No. 1689, European Central Bank.

Ghalayini, L. (2011): ”Does Dollar Weakness Cause High Oil Prices?”, *International Research Journal of Finance and Economics,* 70, 81-89.

Giacomini, R, H. White (2006): ”Tests of Conditional Predictive Ability”, *Econometria,* 74(6), 1545-1578.

Habib, M. M. And M. M. Kalamova (2007): ”Are there oil currencies? The real exchange rate of oil exporting countries”, Working Paper Series No. 839, European Central Bank.

Haug, A. K. (2014) *Oljepris utsetter renteøkning*. Available from: http://www.dn.no/nyheter/politikkSamfunn/2014/10/23/2157/Renten/oljepris-utsetter-rentekning (Viewed: 7 May 2015).

Intercontinental Exchange (2013): *ICE Brent Crude Oil – Frequently-Asked Questions*. Available from: https://www.theice.com/publicdocs/futures/ICE_Brent_FAQ.pdf (Viewed: 7 May 2015).

Jensen, M. C. (1967): ”The Performance of Mutual Funds In the Period 1945-1964”, *Journal of Finance,* 23(2), 389-416.

Mark, N. C. (1995): ”Exchange Rates and Fundamentals: Evidence on Long-Horizon Predictability”, *The American Economic Review,* 85, 201-218.

Melvin, M., J. Prins and D. Shand (2013): ”Forecasting Exchange Rates: An Investor Perspective”, CESifo Working Paper No. 4238, The international platform of Ludwig-Maximilians University´s Center for Economic Studies and the Ifo Institute.

Meese R., (1990): ”Currency Fluctuations in the Post-Bretton Woods Era”, *Journal of Economic Perspective,* 4(1), 117-134.

Mussa, M. L., (1990): ”Exchange Rates in Theory and in Reality”, *Essays in International Finance,* No. 179*,* Department of Economics, Princeton University.

Meese R., K. S. Rogoff (1983a): ”Exchange rate models of the seventies. Do they fit out of sample?”, *Journal of International Economics,* 14, 3-24.

Patell, J. M. and M. A. Wolfson (1984): ”The Intraday Speed of Adjustement of Stock Prices to Earnings and Dividend Announcements”, *Journal of Financial Economics,* 13*,* 223-252.

Rossi, B. (2013): ”Exchange Rate Predictability”, *Journal of Economic Literature,* 51(4)*,* 1063-1119.

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

Sharpe, W. F. (1994): ”The Sharpe Ratio”, *The Journal of Portfolio Management,* 21(1)*,* 49-58.

The Norwegian Petroleum Sector (2014): *Facts 2014*. Available from: http://www.npd.no/Global/Engelsk/3-Publications/Facts/Facts2014/Facts_2014_nett_.pdf (Viewed 7 May 2015).

U.S. Energy Information Administration (2012) *This week in Petroleum September 12*. Available from: https://www.cmegroup.com/education/files/ed172-today-in-petroleum-2012-09-12.pdf (Viewed 7 May 2015).

U.S. Energy Information Administration (2014) *Benchmarks play and important role in pricing crude oil.* Available from: http://www.eia.gov/todayinenergy/detail.cfm?id=18571 (Viewed 7 May 2015).