WISEflow
Europe/Oslo(CEST)
03 Jul 2023

Handelshøyskolen BI

GRA 19703 Master Thesis

Thesis Master of Science 100% - W

Predefinert informasjon

Startdato:	09-01-2023 09:00 CET	**Termin:**	202310
Sluttdato:	03-07-2023 12:00 CEST	**Vurderingsform:**	Norsk 6-trinns skala (A-F)
Eksamensform:	T		
Flowkode:	202310||11184||IN00||W||T		
Intern sensor:	(Anonymisert)		

Deltaker

Navn:	**Felix Skarboe Rønningen**

Informasjon fra deltaker

Tittel *:	Can you use oil and natural gas prices to forecast the norwegian krone?
Navn på veileder *:	Leif Anders Thorsrud

Inneholder besvarelsen konfidensielt materiale?:	Nei	**Kan besvarelsen offentliggjøres?:**	Ja

Gruppe

Gruppenavn:	(Anonymisert)
Gruppenummer:	338
Andre medlemmer i gruppen:	Deltakeren har innlevert i en enkeltmannsgruppe

CAN YOU USE OIL AND NATURAL GAS PRICES TO FORECAST THE NORWEGIAN KRONE?

SUPERVISOR: LEIF ANDERS THORSRUD

PROGRAMME: MASTER OF SCIENCE IN APPLIED ECONOMICS

Preface

This master thesis in applied economics marks the end of a 5-year study run which has been challenging, but also highly rewarding and interesting. The same can be said for my work with this master thesis. I would like to express gratitude towards my supervisor Leif Anders Thorsrud for help with the thesis process. I would also like to thank my family and my girlfriend for their continuous support throughout my educational process.

ABSTRACT

This thesis studies if we can use a commodity prices, namely oil and gas to forecast the NOK/USD exchange using data at a monthly frequency. We analyze the commodities predictive ability through two main models. The first model uses contemporaneous commodity prices as a predictor, while the second uses lagged commodity prices. We use these models to create out of sample forecast one step ahead. Our results are compared to a random walk benchmark model using the Diebold and Mariano test statistic. We find little evidence of oil and gas prices being able to forecast the exchange rate. Only the model with contemporaneous oil prices produced significantly better results than the random walk benchmark.

Contents

1.0 Introduction

This master thesis will explore whether or not oil and gas prices can forecast in the NOK/USD exchange rate. specifically, we'll be using brent oil prices and natural gas prices on a monthly frequency. The research question we will be working with in this thesis is therefore: "Can we use oil and gas prices to predict the Norwegian krone. Oil and gas production and exportation is probably the best-known reason for the solid economic growth Norway has experienced over the years. It is definitely reasonable to say that Norway's economy is heavily dependent on the price of these internationally traded commodities. In 2014, however, something interesting happened. Over night the oil prices plummeted and they continued to drop for quite some time, at a certain point oil price was half of what it used to be. As this happened we also experienced a sharp depreciation in the Norwegian krone against larger international currencies such as the American Dollar, the British Pound and the Euro. This leads one to believe that the price of a nations largest exported commodity is closely tied to the exchange rate, in Norway's case that would be the price of oil and natural gas.

The variables in question in this thesis differs a lot in terms of its impact on the world economy. Oil price and to a degree natural gas price are extremely important variables in the world economy. Both are usually denominated in American dollars and is therefore also usually bought and payed for with this currency. Its therefore highly centralized. This is in stark contrast to the Norwegian krone. When we take into consideration the massive size of the foreign exchange market, its fair to state that the NOK plays a relatively small and insignificant role. It would therefore be unreasonable to expect the Norwegian Krone to have any noticeable effect on the either the global economy or oil/gas prices. More likely, however, is the opposite scenario. That is the scenario that changes in oil/gas prices will impact the Norwegian krone. Something that is worth mentioning is the fact that oil prices has begun to rise again since the sudden fall in 2014, the Norwegian krone however is still continuing its downward spiral. Figure 1 below graphs the oil and exchange rate and the gas price and the exchange rate. When it comes to the oil price and exchange rate the graph suggests that there might be some correlation in the movement between the exchange rate and the oil price. The correlation between them seems to be negative, ie. If the oil price goes up then the Norwegian krone appreciates against the dollar. This is not true for the latest part of the sample though, here the oil price increases sharply, but the krone depreciates against the dollar. A similar relationship can be seen between the gas price and the NOK/USD exchange rate in some parts of the sample, but it seems to be weaker than what is the case with oil prices.

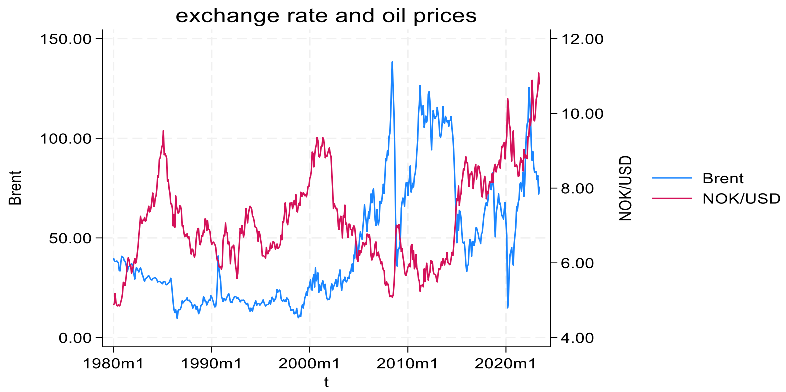

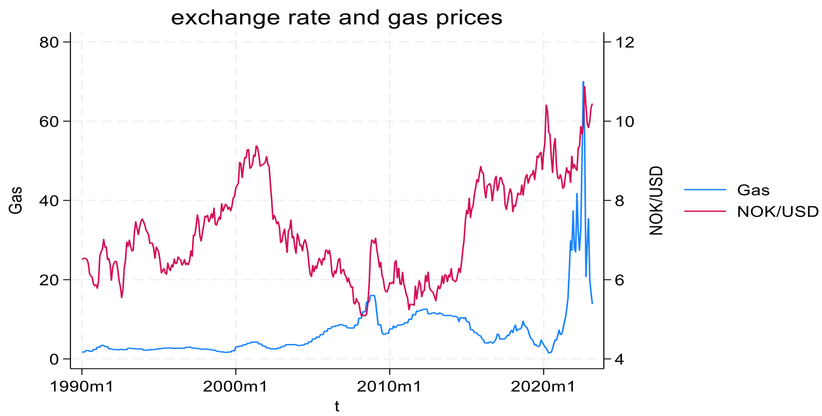

Figure 1 graphing the oil and gas prices together with the exchange rate

This thesis is an extension of a Ferraro et.al(2015) paper that explores the idea of oil price changes being an explanatory variable in changes in exchange rates for nations that heavily depends on oil, either producing on a large scale or heavy importers of it. The paper proposes two main models to explain this, a contemporaneous model and a “true forecasting model”. The contemporaneous model is a normal OLS regression using first differenced oil prices as an independent variable and the NOK/USD as the dependent variable. The true forecasting model instead uses lagged values of oil prices and exchange rates. The authors use these models to forecast the exchange rates. The outcome is compared to a benchmark model which in this case is a random walk model. To compare the two models, the Diebold/Mariano test statistic is employed(1995). Evidence provided by the authors state that they were successful in beating the random walk model when forecasting with the contemporaneous oil price model, when using lagged values for the oil price, results were not as conclusive.

2.0 previous research

Exchange rate movements and fluctuations have been more or less an enigma for economists for a long time, and trying to capture exactly what factors cause fluctuations has been and still is, a hotly debated topic in international macro and finance research. Many models has been tried and tested in order to predict how exchange rates are going to move in the future. Many, if not all, unfortunately run into to the same problem, that is the random walk outperforms it. In this context, it means that the simple random walk model does a better job of predicting the changes in the exchange rate than more complicated theoretical models.

A very central paper in this debate is that of Rogoff an Meese(1983). In this paper they consider an extensive list of exchange rate models an tests their ability to forecast out of sample. The results are compared to that of a random walk. The results, maybe a bit surprising were that the random walk could forecast the exchange rate just as well as the more complicated models, or in some cases outperforming them entirely. At the time this paper was one of a kind and showed how difficult it was to find models that could outperform the simple random walk model. After this paper, the interest was fully sparked, and now we have a large list of research papers covering this specific topic. MacDonald and Taylor(1994) uses the monetary model of exchange rates and compares it to the random walk. The find that forecast using the monetary model has a lower root-mean squared error than the random walk model, both with drift and without. In this experiment they forecasted for the sterling/US dollar exchange rate. Another study that covers the monetary model of exchange rate is that of Finn(1986). In this study she used the monetary model to forecast the Brithis pound/US Dollar exchange rate 1, 6 and 12 months ahead of time. The results were inconclusive, when using the standard monetary model, she found that it was able to outperform the random walk model. When she instead impose rational expectations on the model, they both performed more or less the same. Chinn and Meese(1985) use four structural exchange rate models to forecast the four bilateral rates of the UK, Canada, Japan and Germany. In this case, they found little success, none of the models were able to forecast better than the random walk. MacDonald and Marsh(1997) considers a new approach to beating the random walk. They use the a simple version of the classical purchasing power parity theory to produce long run relations between three bilateral exchange rates, the German mark, the British pound and the Japanese yen. In order to make out of sample forecast they use the long run relations, to make short run dynamic models. These short run models are then used to forecast. In the authors own words, they were able to significantly beat the random walk model in a time span as short as three months.

Moving on to studies that that analyze oil prices and exchange rates, we have Akram (2004). Akram explores a possible relationship between the oil price and the Norwegian krone and reports some interesting findings. Amongst them is the discovery of a strong non-linear relationship between the krone and the oil price. Interestingly, this relationship is at its strongest when the oil price is low. Low oil prices are defined as 14USD per barrel or lower in the specific dataset he was using. Additionally, he discovered that the relationship between the two is stronger when the oil price is on a downward trend. Moreover he uses this negative relationship between the variables to create a model for the exchange rate to be used in a forecasting experiment. The model turned out to be better at forecasting the exchange rate when compared to a model with a linear relationship. More importantly it was also able to beat the random walk benchmark!

Another interesting study is that of Chen et al.(2010). Instead of “just” studying the effects of commodity prices on exchange rates, the also investigate the reverse, meaning they also check if changes in exchange rates have an effect on commodity prices. This study uses commodity price indexes that are specific to the country of interest. All the experiments in this study is carried out using quarterly data. The study concludes that exchange rates can successfully be used to forecast commodity prices, at least compared to the random walk benchmark. The reverse project, however, was not as successful, when they used commodity prices to forecast exchange rates they were not able to beat the random walk benchmark model.

The most important study in the context of this thesis is that of Ferraro et al.(2015). This study covers multiple commodities and exchange rates, the aim was to find out whether or not a countries largest exported commodity could forecast changes in that nations exchange rate. one sub experiment of this paper uses changes in the WTI oil price index to forecast the Canadian dollar and the Norwegian krone, both in terms of the American dollar. Throughout this whole study, daily data is being used. The authors considers two different models. The first one is a model where they employ contemporaneous changes in the WTI oil price, they authors dub this an “out of fit” exercise. This model was a success, it was able to beat the random walk in terms of predicting the exchange rate. The second model they use is more of a “true” forecasting exercise, meaning that they use lagged value for the changes in oil price. Also this model proved to be a success, but not to the same degree as the contemporaneous model. The lagged model had less significant results in addition to being assorted, assorted meaning they could not prove forecasting ability in the entire sample, only parts of it. Nevertheless, this paper was a fairly significant one due to the fact that it was able prove the ability of commodity prices to forecast exchange rates in a short horizon timespan while also being statistically significant.

3.0 methodology

In this section we will be covering what statistical analysis we will performing on our data in order to find out whether or not we can prove if oil and gas prices truly can predict the exchange rate. Before we cover the forecasting experiment, we will discuss some important statistical properties that needs to be present in order for us to be able to do the forecasting experiment.

3.1 Stationarity

Since our forecasting experiment will be carried out using timeseries data, we need to have a grasp on whether or not our timeseries are stationary, otherwise the next steps in our analysis could be compromised. Generally, we say that a timeseries is stationary if the three following properties are present.

1. $E(y_t) = \mu$ (constant mean)

2. $var(y_t) = E(y_t - \mu)^2 = \sigma^2$ (constant variance)

3. $\Upsilon_s = [(Y_t - \mu)(Y_{t-s} - \mu)]^2$ (constant autocovariance)

If all these criteria are met, we have what we call a strictly stationary process, a weak stationary process only requires the mean and the variance to be constant. This allows the timeseries to have a drift or a constant. We say a time-series is stationary if it is mean-reverting, meaning a process that returns to certain mean over time. If the time-series we are working with is not stationary it could cause some grave consequences to our analysis. If we were to use a non-stationary time series in an OLS regression, we would mostly likely end up with a so-called spurious regression. Amongst other things, this would lead to explosive estimates which means the estimates our model makes moves further and further towards infinity, the further ahead we estimate. This is obviously not a desirable quality for our model to have, since our model would then yield an unrealistic view of reality. Another serious consequence on faces with spurious regressions is that the standard errors the model yields would be invalid. This means we lose the ability to perform t and f-tests on our model in normal ways. In our case, the Diebold-Mariano test would be invalid.

3.2 unit root

Another way a time-series can be non-stationary other than breaching the criteria we listed above, if it contains a unit root. We say that a time-series has a unit root if the absolute value of the time-series coefficient is larger or equal to zero. A well known example of a time series with a unit walk is the random walk with drift:

$$Y_t = \alpha + \phi Y_t \ + \epsilon_t$$

Where:

Yt = the variable Y at a given time t

α = trend parameter

φ = coefficient for the lagged value

εt = zero-mean white noise error term.

If we have that φ>1 we have that the time-series has a unit root. In this case shocks will lead to permanent changes if we were to use the time-series in a model. Luckily for us, there is a method to get rid of this problem that works in most cases, namely differentiating the time-series. This involves taking the first-difference of the series and substituting in Yt-1. With the random walk model we used as an example earlier, first-differencing would yield the following result:

$$\Delta Y_{t-1} = \alpha + \epsilon_t.$$

This process is stationary and could be safely be used in an OLS-regression. This ties us nicely into the term integration. We say that a time-series that needs to differenced once in order to be stationary is integrated of the first order or *I*(1), a series that needs to be differenced twice in order to be stationary is integrated of second, *I*(2) and so on. If the series is already stationary without us having to differentiate it, it is *I*(0).

3.3 The Augmented Dickey-Fuller test

In order for us check whether or not the time series data we will utilize in this thesis contains a unit root or not, we will be using the Augmented Dickey fuller test (ADF-test). The ADF-test is an expansion of the original Dickey-Fuller test (DF-test) which was developed in a 1979 paper. The mechanisms are quite similar in both so we will start by explaining the basics of the DF-test before we move on to the ADF-test. Our starting point is the following simple AR(1) process:

$$Y_t = \phi Y_t + u_t$$

As we remember from before, Yt has an unit root if ϕ is larger or equal to 1. If that is the case Yt will be non-stationary. Since this time-series is only a sample, we will never know the true value of ϕ. Instead, we estimate it and form the following hypotheses.

H0: $\phi = 1$

HA: $\phi < 1$

In practice we first-differentiate the AR(1) process so we have:

$$Y_t - Y_{t-1} = \phi Y_{t-1} + u_t - Y_{t-1}$$

- ⇨ $\Delta Y_t = (\phi - 1)Y_{t-1} + u_t$
- ⇨ $\Delta Y_{t-1} = \psi Y_{t-1} + u_t$

Where $\psi = 1 - \phi$. If Yt has an unit root which is the case if ϕ=1, then Ψ =0. The hypotheses in the DF-test will therefore be:

H_0: $\psi = 0$ (the time series has an unit root)

H_A: $\psi < 0$

The test-statistic for the DF-test is given by $\tau = \frac{\psi}{SE(\psi)}$. As one might notice, this test statistic heavily resembles that of the test-statistic for a normal t-test. Although they are quite similar, we cannot use critical values from the t-distribution, instead Dickey and Fuller created a new set of critical values suited for this test.

We have now shortly described the DF-test and can move on to how it differs from the ADf-test that we will be using. The main difference between the two is that the ADF can be used for AR(p) models, while the DF-test can only be used for AR(1) models. Here p is the number of lags included. The reason behind the extension up to p lags is that the DF-test assumes that the error term is white noise which does not necessarily have to be the case. For the ADF-test we therefore have the following process.

$$\Delta Y_t = \psi Y_{t-1} + a_1 \Delta Yt - 1 + a_2 \Delta Y_{t-2} + \cdots + a_p \Delta Y_{t-p} + u_t$$

$$\Rightarrow \ \psi Y_{t-1} + \sum_{i=1}^{p} a_i \Delta Y t_{t-1} + \ u_t$$

3.4 Our main models

This section develops and explains the models we will estimate in this thesis and the models we will use to forecast the exchange rate.

The study of Ferraro et.al(2015), which we will adopt uses the following two main econometric specifications throughout the study

$$(1) \quad \Delta s_t \ = \alpha \ + \beta \Delta p_t + u_t \ \Delta, t = 1, \ldots, T$$

$$(2) \quad \Delta s_t \ = \alpha \ + \beta \Delta p_{t-1} + u_t \ \Delta, t = 1, \ldots, T$$

Model (1) is the model that uses contemporaneous commodity prices, and (2) is the model that uses lagged commodity prices. St is the natural logarithm of the exchange rate, pt refers to the natural log of the commodity price in question, in our case oil and gas prices. The delta symbol means the variable behind it is first differenced. T refers to the model sample size.

These models are the base for the models we are going to use for forecasting, with that in mind, the forecast for the contemporaneous model is therefore:

$$\widehat{s_{t+1}}^{f} = \widehat{\alpha_t} + \widehat{\beta_{t+1}} \Delta p_t \ \ , t = R, R+1, \ldots., T-1 \qquad (3)$$

The α and β in this case refers to parameter estimates that we obtain from a rolling sample of observation, R is the sample estimation window. Ferraro et al. states that model (3) actually does not give us a traditional out of sample forecast since we do not used lagged values. Instead we use changes in oil/gas price to forecast changes in exchange rate, but both of these changes are being recorded on the same day. Instead of a true forecasting exercise we therefore have an out of sample exercise. We therefore have that $\widehat{s_{t+1}}^{f}$ refers to the one step ahead out of sample fit of the exchange rate tomorrow.

For the true forecasting exercise that utilizes lagged commodity prices, we have the following model:

$$\widehat{s_{t+1}}^{f} = \widehat{\alpha_t} + \widehat{\beta_{t+1}}\Delta p_t \ , t = R, R+1, \dots, T-1 \qquad (4)$$

Again, we have that α,β refers to parameters that we obtain from a rolling sample of observation, the R refers to the in-sample estimation window. $\widehat{s_{t+1}}^{f}$ will be the one step ahead out of sample forecast for the exchange rate.

3.5 The Diebold-Mariano test

In order for us to find out if our models can reliably beat the random walk model, we need some sort of test. The Diebold Mariano test is what we will use. Originally this test was developed in a 1995 paper by Francis Diebold and Roberto Mariano. This test was invented as a way to formally decide if one model produces significantly better forecast results than another model. Essentially how this test works is that it calculates the difference between two models mean-squared forecast errors and finds if there is a statistically significant distance between the two models. The intuition behind the test is the following:

Define errors for the forecast experiment as follows.:

$$e_{it} = \hat{y}_{it} - y_t, \ \ t = 1{,}2$$

We assume that loss associated with forecast error i, is a function of the forecast error e_{it}. We denote this function as $g(e_{it})$. This function g(.) is a loss function and has the following properties:

1. Is zero when there is no error made.
2. Can never be negative.
3. Becomes larger, when errors increase in magnitude.

From this we can define loss differential between our two forecasts as:

$$d_t = g(e_{1t}) - g(e_{2t})$$

We say that the forecast has equal predictive accuracy if and only if the loss differential has zero expectations for all t. From this we can form the following null and alternative hypotheses.

$$\boldsymbol{H_0: E(d_t) = 0}$$

$$\boldsymbol{H_A: E(d_t) \neq 0}$$

The null hypothesis states that there is no difference in forecast accuracy between the two models, the alternative hypothesis states that the two models produce forecast with a varying level of accuracy.

We that under the null hypothesis, the Diebold-Mariono(DM) test statistic is distributed asymptotically with N(0,1). Therefore, we have that the null hypothesis will be rejected if the DM-statistic we calculate is outside the following range:

$$|DM| > Z_{\alpha/2}.$$

The $Z_{\alpha/2}$ value is the positive critical value we get from the standard normal distribution table when we use half of the significance level α. We use half of the significance level since this is a two-sided test.

3.6 Steps

This section shortly describes how and in what order methodology is uses in order to answer the research question in our thesis.

Step 1

The first thing we need to find out after loading our data into the software and transforming them, is to find out if they are stationary or not. To do this we use the ADF-test described in section 3.3. We test for this by using the ADF-test statistic and compare them to critical values from the interpolated Dickey-Fuller table. If the test concludes that the variables are stationary we can safely move on to next step.

Step 2

Since we have that all our variables are stationary we can safely use OLS regression to estimate model (1) and (2). We can also use models (3) and (4) to produce forecasts for the exchange rate using oil and gas prices as the predictor respectively.

Step 3

We store results from the forecast experiment in step 2 and use the results to calculate the DM-statistic. This test statistic is then compared to critical values from the normal distribution table at the following levels of significance, 1%,5%,10%.

4.0 Data

In this section we describe the datasets we use in this thesis, we graph them and explain in detail why and what transformations we perform on the data in order to do our forecasting experiment.

4.1 Data sets

This thesis uses two main dataset in order to answer the research questions, both varying in length and frequency. The first dataset contains NOK/USD exchange and spot brent oil prices. For clarification, the exchange rate is denominated in how much Norwegian krone you need to purchase one American dollar. This dataset is in monthly frequency. The other data set we use is one containing the NOK/USD exchange rate and Global price of natural gas, EU. This dataset is monthly in frequency. The two datasets vary a bit in length as well, the oil price set contains data from January 1980 to March 2023. The gas price set on the other hand contains data from January 1990 to March 2023. The oil price dataset was retrieved from global financial data, while the natural gas price dataset was retrieved from FRED. The measurements for the commodity prices are as follows: oil is measured in USD per barrel, while the gas price is measured in USD per cubic feet.From this point onwards, when we talk about exchange rates, we mean the NOK/USD exchange rate.

4.2 Why did we chose the Brent oil in our analysis:

Even though this thesis is based on the work of Ferraro et.al(2015) which uses WTI oil prices, we choose instead to use the Brent oil price. We recon that there is three main types of oil sold on the large international oil market. Brent oil, which is produced in the North sea, the West Texas intermediate (the WTI), this is usually considered as the main benchmark for oil in USA. Lastly we have the Dubai/Oman oil which largely dominates the Asian market. There is mainly one reason why I have chosen to work with the Brent oil in this paper is that brent oil is the type of Oil that Norway extracts, and exports internationally. That being said both WTI and Brent oil are quite similar types of oil and their price movements are also quite similar.

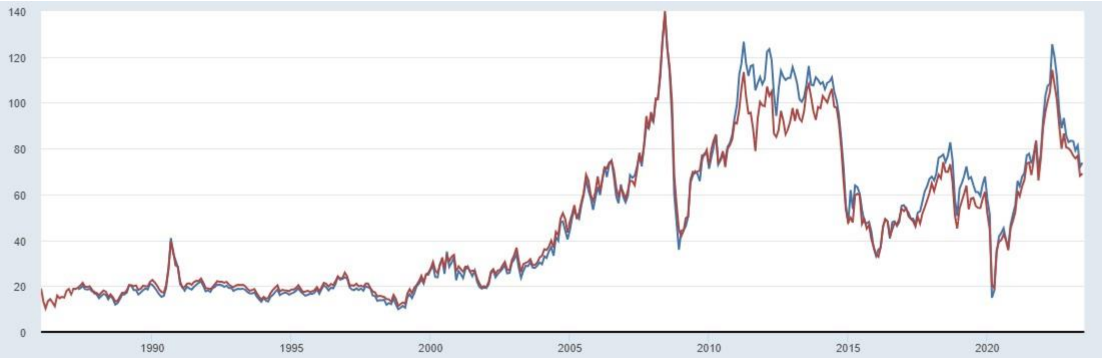

Figure 2: graphs Brent and WTI oil prices, Brent is the blue line and WTI is the red

From figure 1 we can see that they move quite similarly, this graph also indicates a high positive correlation between the two. The trend seems to be that WTI used to be priced a little bit higher, at least up until about 2010, when the Brent oil prices higher. This period is also the time where the difference between the two oil prices is at its greatest.

4.3 exchange rate

In figure 2 we observe the NOK/USD time-series with a daily frequency. As we remember we defined this as how much Norwegian krone you need to purchase one American dollar. If we observe a high value in the graph, it therefore means we have a weak krone, and a low value we therefore have a strong krone. Using the same logic, a downwards trend in the graph means the krone is appreciating against the dollar, and a rising trend depicts a depreciation of the krone against the dollar.

Figure 3 The NOK/USD exchange rate

The graph as a whole does not seem to have any particular trends. The krone was at its strongest in years before the financial crisis in 2008, right after though it depreciated strongly, before slowly apricating again as the financial crisis started to die out. Since the oil price shock in 2014 up until this point in time, the Norwegian krone has depreciated strongly against the dollar. This is no small part due to three massive moments for the world economy. The mentioned oil price shock, the Corona virus outbreak and lately Russia's attack on Ukraine. During crises like these, investors in the foreign exchange market tends to stay clear of smaller currencies such as the Krone, in favor of larger currencies such as Dollars, pounds etc. This may explain some of the depreciation in the krone, but there are probably several other factors in play as well.

4.4 Gas

In figure 2 we display movements in the price of natural gas from the start of our sample period in 1990 to today. Throughout most of sample, price movements do not seem to be too volatile, that is until recent years. In the last part of the sample we see a massive increase in price, before being followed by an almost equally large decline. The massive increase in prices is likely due to the war in Ukraine changing the underlying dynamics in the gas market. Most of this comes down to Europe being affected by large reductions in Russian production and exports of natural gas. In fact, Russia cut its exports down to 53% of what it was on average pre-war in the period 2017-2021. (Agnolucci, Nagle, Temaj, Worldbank)

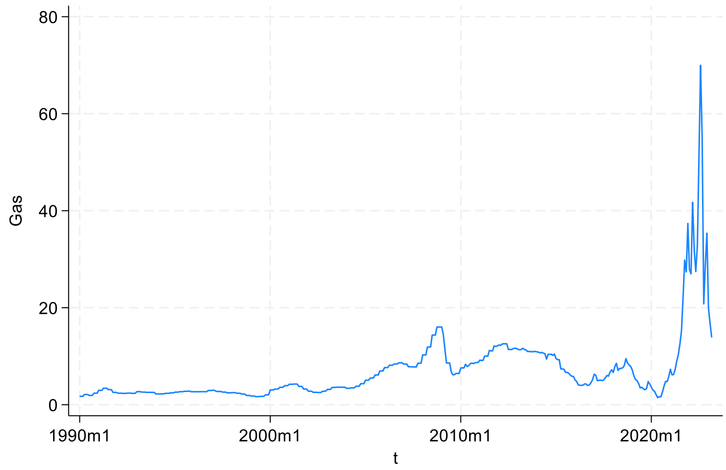

Figure 4 Gas prices from 1980-2023

4.5 The transformed data

In the main analysis in this thesis we use the first differenced natural logarithm. This proves useful in many ways. By using the natural logarithm of a variable as we obtain the growth rate of the variable by doing so. Logarithmic dataset also help by reducing the risk of possible heteroscedasticity and also mitigates possible extreme values. This in turn, mitigates the risk of breaching the classical assumptions a linear model needs to fulfill. We first difference the variable to avoid possible non-stationarity in the data, although we also formally test just to be absolutely sure we a working with a stationary series. Figure 5 shows the first differenced natural logarithm of our variables, as we can see they seem do not seem to have any trends, they also cross their mean value frequently. This indicates that our variables are $I(1)$, but we will still formally test just to be absolutely sure.

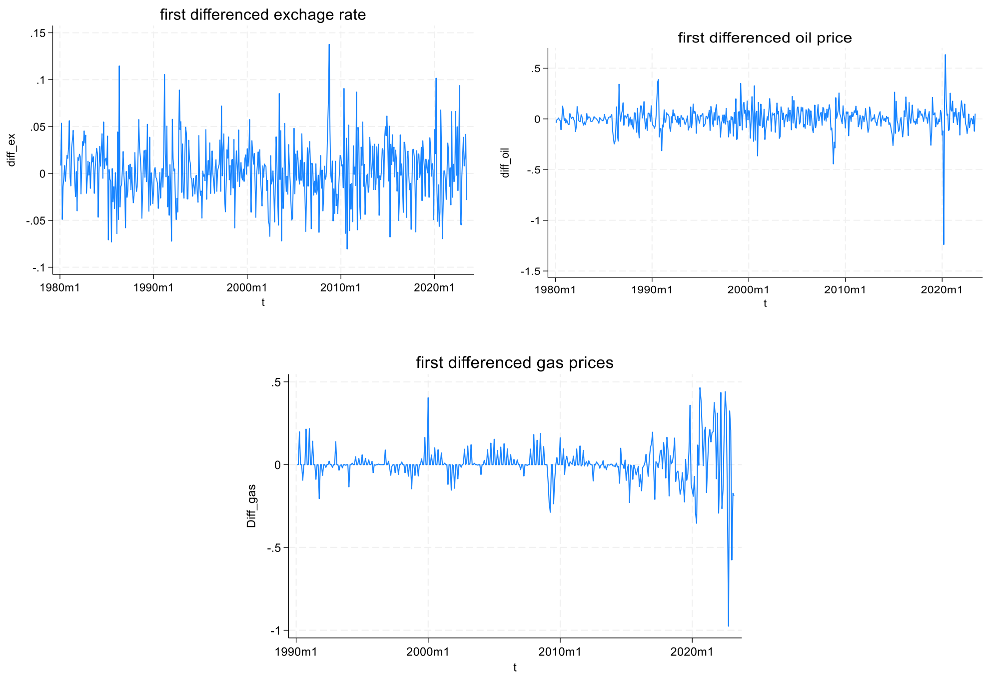

Figure 5 First differenced variables

5.0 Empirical results.

In this section, we perform our forecasting experiment as described in section 3.4 and use the Diebold-Mariano test to compare it to the random walk model. We present and analyze results from this experiment, therefore presenting a conclusion on whether or not we can use commodity prices to forecast exchange rates.

5.1 ADF-test

The very first thing we do, is check for unit roots in our time-series data, thereby finding out if our data is stationary or not. To do this we deploy the ADF-test described in section 3.3. we only test for unit roots in the transformed data since these are the data we will be using, and the non-transformed data is unlikely to be stationary anyway. The results from the test are posted in table 1. As we can see, all the ADF-test statistics are outside the range of the critical values from 1%-10%. We can therefore conclude that none of our first-differenced variables contains an unit root. We can therefore safely use them in an OLS regression.

Variable	ADF-test statistic	1% critical value	5% critical value	10% critical value
Exchange rate	-15,6283	-2,584	-1,957	-1,631
Oil price	-14,5268	-2,584	-1,957	-1,631
Gas price	-14,579369	-2,587	-1,939	-1,632

Table 1 ADF test results

5.2 estimating the exchange rate

The analysis in this paper is solely based on a simple model where our commodity prices are the only variables that explain changes in the exchange rate. Below you will find results estimation results from models (1) and (2) where we use both the oil price and the natural gas price to explain changes in the exchange rate. As we remember, model (1) is the model that uses contemporaneous commodity prices, and model (2) is the model that utilizes lagged commodity prices. All our parameters will be estimated using OLS.

5.2.1 the contemporaneous model

	Estimate	SE	tStat	pValue
(Intercept)	0.012557	0.0093111	1.3486	0.17805
x1	-0.017854	0.0018012	-9.9122	2.5114e-21

Figure 6 regression results with contemporaneous oil prices

	Estimate	SE	tStat	pValue
(Intercept)	0.0011908	0.0016223	0.73401	0.46338
x1	-0.00079977	0.013847	-0.057757	0.95397

Figure 7 regression results with contemporaneous gas prices

As we can see from figure 1, we estimate the value of our oil price parameter to be -0,01785. Since we are using logarithmic values, this suggests that if the oil price were to increase by 1%, we should expect the exchange rate to decrease by 0,01785 percentage points on average. We estimate the intercept to be 0,01256, it is however insignificant on 1%,5% and 10% levels, suggesting that the exchange rate does not change when we have no change in the oil price. As for the model with gas prices, we can see that that contemporaneous gas prices has a very small and statistically insignificant effect on exchange rates.

5.2.2 the lagged model

	Estimate	SE	tStat	pValue
(Intercept)	0.011532	0.010128	1.1385	0.25542
x1	-0.0042107	0.0019584	-2.15	0.032018

Figure 8 regression results with lagged oil prices

	Estimate	SE	tStat	pValue
(Intercept)	0.0010076	0.0016167	0.62328	0.53346
x1	0.03075	0.013826	2.2241	0.026707

Figure 9 regression results with lagged gas prices

When we use lagged oil prices instead of contemporaneous prices, we can see that the estimated relationship is much weaker. A one percent increase in lagged oil price growth rate will only lead to a 0,00421 percentage points drop in the exchange rate. This time, the estimated coefficient is also no longer significant on the 1% level, it is however significant when using 5% and 10% significance levels. Again we have that the intercept is not statistically significant. The model with lagged gas prices breaks the trend in several ways, firstly this model estimates a stronger relationship with lagged values, this is the other way around for oil prices. Secondly it is the only model that suggests a positive relationship between commodity prices and exchange rates. We can see that our model suggests a 1% change in the lagged gas prices, should result in a 0,03075-percentage point change in the exchange rate. this relationship is statistically significant on the 5% and 10% levels as well as we see from the p-value. Again we have that the intercept is non-significant.

5.3 Doing the forecasting experiment

In this section, we will answer the research question on whether or not our chosen commodity prices can predict the exchange rate. We adapt a similar approach to that of Ferraro et.al which utilizes a contemporaneous out of sample fit exercise and true forecasting experiment using lagged commodity prices. We extend upon their research both by using a more recent dataset and doing the experiment with natural gas prices. Instead of using the WTI oil prices as Ferraro et. al do, we use Brent oil prices since it is more specific to Norway. The results of the experiment is not entirely conclusive, we only find evidence of oil prices being able to predict the exchange rate better than the random walk when we use the model with contemporaneous prices. The model with lagged oil prices could not perform the random walk on any significance levels. We found the same result with gas prices, with both contemporaneous and lagged values.

5.4 Results from the contemporaneous model

We use model (3) that we discussed in in section 3.4. We use this model to estimate our parameters via a rolling OLS regression. This exercise will use realized changes of our commodity prices to determine how good our model is at predicting exchange rates, it is therefore not a true forecasting exercise. That being said, if we find success with this experiment, an also find a good model to predict future oil prices, we then have a better chance of also making good estimates of future exchange rates Ferraro et.al(2015). We perform this experiment using multiple in-sample window sizes.

5.4.1 results oil prices

Table 2 shows the result from the model with contemporaneous oil prices. If we report a Diebold Marioano test-statistic below that of -1,645, we have found proof that the model with contemporaneous price significantly outperforms the random walk at 10% significance level. As you can see from our negative DM-test statistics, we were able to outperform the random walk for all selected window sizes. Window sizes 1/2 through 1/5 were able to beat the random walk at 1% significance, while the rest of the window sizes outperformed it at the 5% significance level.

	1/2	1/4	1/5	1/10	1/20
DM test stat	**-3,658**	**-2,934**	**- 2,781**	**-2,285**	**-2,056**
Crit value α=10%	-1,645	1,645	1,645	1,645	1,645
Crit value α=5%	-1.96	-1.96	-1.96	-1.96	-1.96
Crit value α=1%	-2,33	2,33	2,33	2,33	2,33

Table 2: test results from the model with contemporaneous oil prices

5.4.2 results gas prices

When using the model with contemporaneous gas prices, we were not successful in producing significantly better results than the random walk. We had to keep the null hypothesis of no significantly different forecast accuracy for all significance levels. Results from this experiment are posted in table 3 below.

	1/2	1/4	1/5	1/10	1/20
DM test stat	**-1,231**	**-0,873**	**-0,937**	**-0,852**	**-0,679**
Crit value α=10%	-1,645	1,645	1,645	1,645	1,645
Crit value α=5%	-1.96	-1.96	-1.96	-1.96	-1.96
Crit value α=1%	-2,33	2,33	2,33	2,33	2,33

Table 3: test results from the model with contemporaneous gas prices

5.5 results from the model with lagged prices

Having covered the model with contemporaneous commodity prices, we now use the true forecasting model with lagged commodity prices. For this experiment we use model (4) from section 3.4. Again, we estimate the coefficients using OLS with a rolling sample window of observations. Finally we compare our results to our random walk benchmark model using the DM-test. We found little success with this experiment; our forecast models were not able to significantly outperform the random walk at any significance level, this held true for all window sizes.

5.5.1 results lagged oil prices

Table 4 depicts results from the Diebold and Mariano test with lagged oil prices, as we can see we failed to outperform the random walk on all significance levels. This holds true for all window sizes used for the test as well.

	1/2	1/4	1/5	1/10	1/20
DM test stat	**-1,186**	**-0,916**	**-0,907**	**-0,583**	**-0,349**
Crit value α=10%	-1,645	1,645	1,645	1,645	1,645
Crit value α=5%	-1.96	-1.96	-1.96	-1.96	-1.96
Crit value α=1%	-2,33	2,33	2,33	2,33	2,33

Table 4: test results from the model with lagged oil prices

5.5.2 results lagged gas prices

table 5 shows the results of the Diebold-Mariano test when we use the forecasting model with lagged gas prices. Similarly to the model with lagged oil prices, we fail to beat the Random walk benchmark on any significance level, again this is true for all window sizes.

	1/2	1/4	1/5	1/10	1/20
DM test stat	**-0,946**	**-0,739**	**-0,852**	**-0,647**	**-0,428**
Crit value α=10%	-1,645	1,645	1,645	1,645	1,645
Crit value α=5%	-1.96	-1.96	-1.96	-1.96	-1.96
Crit value α=1%	-2,33	2,33	2,33	2,33	2,33

Table 5: test results from the model with lagged gas prices

6.0 Limitations

There are many different angles you could look at this study from, we could not explore all of them in this thesis. We therefore believe that there is much more research that can be done on this topic. First and foremost ,due to many reasons, this thesis is limited to monthly data. The data we use in this thesis are considered as financial assets. Participators in these market will have access to data with very high frequency. The will utilize this to find price discrepancies in the market to make good trades. For this reason it may be the case that monthly data is at to low a frequency to reflection information that could be relevant to forecasting the exchange rate. one could therefore suggest using data at a higher frequency such as weekly, daily or maybe even hourly data. Another factor this thesis does not consider is the fact that during major parts of our sample the Norwegian Krone was not a free floating currency as it is today, rather it was pegged to a number of different currencies. We have not compensated for that in any way in this analysis. In our time period a lot of political affairs could have effected swings supply and demand of our variables as well, particularly the oil and gas. This includes a larger focus on renewable energy sources and new climate regulations, Russia's war against Ukraine and so on. There is definitely a chance these political affairs could have produced swings in the oil price which may in turn have yielded more explanatory power on the exchange rate. This not taken into account in this thesis, but might be interesting to look into in another study.

7.0 Conclusion

In this thesis we set out to find whether or not we could use oil and gas prices to forecast the NOK/USD exchange rate. To answer this question we adopted the methodology of Ferraro et.al(2015) which utilizes a contemporaneous out of sample fit model, and a true forecasting model using lagged commodity prices. We used monthly frequency for the exchange rate, oil and gas prices. We compared our forecast results to that of the random walk model which is said to be the hardest benchmark to beat. To find out if our models could forecast the exchange rate significantly better than the random walk we used the Diebold-Mariano(1995) test. We found little success with this experiment, the only model to outperform the random walk was the model using contemporaneous gas prices. This model however beat the random walk on all significance levels.

8.0 reference list

Akram,Q Farooq. (2004): *"Oil prices and exchange rates: Norwegian evidence*." Econmetrics journal, volume 7, pp. 476-504.

Chenn, Y, K Rogoff and Rossi B (2010): *"Can exchange rates forecast commodity prices)"* The Quarterly Journal Of Economics, August 1145-1194

Chinn, Menzie D and Meese A Richard (1995): *Banking on currency forecasts: How predictable is change in money?* Journal of international economics 38:161-78

Dickey, D and Fuller W (1979): *"Distribution of the estimators for autoregressive time series with a unit root"* Journal Of The American Statistical Association (366a), 427-431.

Diebold, F and Mariano, R (1995): "*Comparing predictive accuracy"* Journal of Business and Economic Statistics" 13, 253-265

Ferraro, D, Rogoff, K and Rossi, B (2015) "*Can oil price forecast exchange rates? An empirical analysis of the relationship between commodity prices and exchange rates."* Journal of international money and finance 54, 116-141.

Finn, Mary G (1986): ""*Forecasting the exchange rate: A monetary or random walk phenomenon?"* Journal of international money and finance, June 1986, 181-220.

Macdonald, R, and Taylor, P,M (1994). "*The monetary model of exchange rate: Long-run relationships, short-run dynamics and how to beat a random walk."* Journal of international money and Finance 13, 276-290.

MacDonanld, R, Marsh, W,I (1997): "*On fundamentals and exchange rates: A Casselian perspective"*. Review of Economics and Statistics 79. 655-664.

Meese, R, Rogoff, K (1983). "*Exchange rate models of the seventies. Do they fit out of sample?"* Journal of international economics 14, 3-24 https://blogs.worldbank.org/opendata/bubble-trouble-whats-behind-highs-and-lows-natural-gas-markets