Norwegian School of Economics
Bergen, Spring 2018

AIS data and the price of Oil

A study of predictive feasability

Aslak Wøllo Flaate and Maria Nikitina

Supervisor: Roar Os Ådland

Master Thesis, ECO and ENE profiles

NORWEGIAN SCHOOL OF ECONOMICS

Acknowledgements

We would like to express our appreciation to Professor Roar Ådland, our thesis supervisor, at whose suggestion this topic was chosen, for this research opportunity and especially for his patience with our time-consuming paper.

Bergen, June 2018

Aslak Wøllo Flaate Maria Nikitina

Abstract

Compared to the oil market, physical movement of oil-carrying vessels is very precise and reflects the real production rates more timely than official reports. In this paper, we examine whether detailed information on crude oil movements, obtained from AIS tracking system, can be used to better predict the oil price. We use a variety of model specifications and introduce a novel instrument for the role of expectations to this question. This instrument is based on vessel speed, and it offers insights into the apparent lack of empirical indications of speed optimization. We show that the AIS data can contribute to predicting the oil price. We also explore Kilian's (2009) hypothesis that the model of the oil price should include three factors: expectations of future prices in addition to supply and demand. We triangulate his instrument with one that we construct independently.

We have explored several different specifications for the relationship between inter- and intraregional oil ship traffic and the oil price and found that a statistically significant relationship exists. Our findings indicate that the correlation between variables make OLS an unsuitable tool for this analysis since endogeneity bias will suppress the actual relationship. However, we have found that the relationship is robust to different VAR specifications. The contribution to explanatory power as measured by Factor Error Variance Decomposition is marginal, but it might still be a small improvement on present methods. We have examined the apparent paradox of non-optimal ship speed behavior. We found that a less stringent specification apparently resolved the paradox; the freight rates do indeed influence ship speed if lags and correlation are allowed. The short time-span is preventing us from conclusively saying that the issue is resolved, but it appears at least to be worthy of further investigation. We assessed the validity of using ship opportunity cost as a measure for GDP. While we cannot address the question of possible bias created by Kilian's use of the Baltic index, we nonetheless offer conceptual support, as our unrelated instrument for the same opportunity cost showed strong statistical significance.

Contents

1. Introduction

Global demand for energy continues to grow as prosperity increases, with oil remaining to be the world's main source of energy, influencing the pricing of the two next largest sources, coal and gas (BP, 2018). Oil is not only a source of energy, but also a dominant cost-driver for everything from freight costs to food production, and a base for derivatives in the financial markets. Because of its great importance, a good understanding of its price behavior has enormous value.

Unfortunately, the oil price is difficult to predict. It exhibits strong heteroscedastic behavior, where it be stable for long periods of time, but can change suddenly if the market believes that some fundamental factor has changed. Oil has very low short-term elasticities, both on the supply and demand side, and consumers of oil-based products often lack good alternatives. To better understand how these factors interact, Kilian (2009) has created a model where the oil price is assumed to be driven by expectations, in addition to the traditional factors of supply and demand.

The intuition behind this idea lies in the combination of the previous points: the oil price can be very unpredictable, and oil consumers are often dependent on oil to be able to generate value. It is therefore reasonable that consumers would worry about future prices and take steps to hedge against unfavorable price changes. The heteroscedastic nature of the oil price makes this costly, however, since the sudden price changes that the oil price exhibits are expensive to insure against. Additionally, the oil price reverts only slowly to its long-term equilibrium, so the one would have to buy coverage for a long period of time to offset the price risks. As a result, buying long-term insurance is an expensive solution. Oil storage is also an expensive solution. Lacking alternative strategies against the uncertainty, the agents in the markets are forced to monitor the market for signs of oil price developments, and therefore the market responds mostly to itself. Kilian (2009) described it as expectation-driven demand. Mathematically this can be described as a strongly persistent series with unstable equilibria can be unstable in the short-run, and a strongly heterogeneous series in the long-run.

In this paper, we examine whether detailed information on crude oil movements, obtained from AIS tracking system, can be used to better predict the oil price. We try different model specifications and use an instrument to control for the role of expectations. The instrument is based on vessel speed, and our findings offer insights into the apparent lack of empirical indications of speed optimization. Specifically, we show that in specifications that are robust to endogeneity, we find highly significant Granger causality between the oil price and tanker speed, in both directions. This indicates that the opportunity costs of ships are indeed a factor that determines speed.

Compared to the oil market, physical movement of oil-carrying vessels is very precise and reflects the real production rates more timely than official reports, which makes it an interesting case-study. Our null hypothesis is that oil market is efficient, which is assumed by economic theory, and that there are no other effects, such as physical oil flow, defined as the volume of oil at a given time transported from one place to another, or combination of flows, that can have a significant effect on the oil price development.

The question of interest is whether we might find any significant results that counter this weak efficiency hypothesis. If so, it means that it is possible to get a better understanding of the oil price path through the inclusion on non-typical data. If not, we will have confirmed that the oil price behaves in accordance with the weak efficiency hypothesis. The contribution of our work is a viability test for AIS data in prediction, where we test whether the data confer significant information. Additionally, we explore whether the vessel speed optimization paradox can be resolved using less restricted specifications. In doing so, we also triangulate Kilian's findings, which have been the subject of legitimate critique. This research could be of interest not only to scholars specializing in oil markets efficiency and shipping, but also to oil traders, and companies involved in oil production and refining.

The remainder of the paper is organized as follows: Section 2 highlights relevant literature, in Section 3 we explain the theory and methodology we use and outlines model specifications, Section 4 describes the empirical data used for this research, Section 5 discusses results. The paper is concluded with Section 6.

2. Literature review

Oil has been a subject of intensive research for a long time. Many papers investigate the crude oil market and prices from different perspectives. The main body of literature focuses on the derivatives market, with the relationship between markets, price discovery, and the spot-futures relationship as central questions; and on physical seaborne trade, with the effect of oil price on vessel speed, freight rates and bunker fuel cost.

Back in 1991, Green and Mork studied the monopolistic behavior of OPEC countries during the time when most of the oil was sold at “official prices” under long-term contracts and only a marginal part of oil trade was done at the spot market. It turned out that the price under the long-term contracts deviated systematically from the ex-post spot price at the time of delivery. However, authors concluded that it was improving over time. Silvapulle and Moosa (1999) examined the lead-lag relationship between oil spot and futures prices of WTI using linear and non-linear models on daily observations from 1985 to 1996. The linear model revealed that futures lead spot prices; however, non-linear model results suggested that effect was in both directions. Ewing and Harter (2000) studied the inter-relationship of UK Brent and US Alaska North Slope oil prices on monthly data from 1974 to 1996. They found that markets were unified and followed a random walk pattern, moving together over time and reacting simultaneously to shocks, with the Alaska North Slope following and adjusting to innovations in the market for Brent.

There are some studies that investigated the degree of crude oil market efficiency over time. Alvarez-Ramirez et al. (2002), Serletis and Andreadis (2004) found evidence of long-range dependence phenomena. Tabak and Cajueiro (2007) extended the studies and focused on how efficiency evolves over time. Their finding was that efficiency does increase and that WTI is more weakly efficient than Brent. However, their findings turned out to be inconsistent with the results of Charles and Darné (2009), who claimed that Brent oil market was more weak-form efficient than the WTI crude oil market in 1994-2008. Both Alvarez-Ramirez et al. (2008) and Elder and Serletis (2008) concluded that markets do show short-term inefficient behavior, yet they become more efficient over a long-term period. Mensi et al. (2012) used Symbolic Time Series and modified Shannon entropy to examine the time-varying degree of informational efficiency of crude oil markets on daily returns for WTI and Brent and supported the conclusion of weak-form efficiency evolution.

Maslyuk and Smyth (2008) used Lagrange multiplier unit root tests with structural breaks to examine the unit root behavior. They found that each of the series can be characterized as random walk process and concluded that, for the time frame considered in the study, it is impossible to forecast future movements in crude oil price based on past behavior.

Another approach was used by Lin and Tamvakis (2001) to study investigated information transmission mechanism between WTI and Brent futures prices in two geographically separated markets, NYMEX (New York Mercantile Exchange) and IPE (International Petroleum Exchange in London). They found out that spillover effects do exist when both markets trade simultaneously. They attributed it to the fact that market participants are flexible in moving from one market to another throughout the trading day and have trading positions at both markets.

Milonas and Henker (2001) focused on the effect of fundamental measures on the oil price. Dées et al. (2007) distinguished between two production behavior models: competitive and cooperative. Ewing and Thompson (2007) studied co-movement of oil price with output, consumer prices, unemployment and stock prices. Kaufmann et al. (2008) incorporated changes in refinery sector into the model and argued that low refinery utilization rate leads to the preference for higher quality crudes and the increase in crude oil prices on average. Kilian (2009) has created a model of the oil price that has three fundamental factors rather than the traditional two. In addition to supply and demand, he claims that expectations of future prices should be a factor by itself.

Kaufmann and Ullman (2009) examined innovations in the oil price and links between spot and futures prices. They find that prices respond to market fundamentals and are exacerbated by speculation. Coleman (2012) expanded the approach by studying fundamental measures but also incorporating proxies for speculative and terrorist activity and dummies for industry events to explain the dynamics of crude oil price.

Alizadeh and Nomikos (2004) studied cost-of-carry relationship. Their hypothesis was that as crude oil is a storable commodity, cost of carry relationship must hold, and physical crudes should be linked to the WTI futures through the transportation costs and tanker freight rates. However, they found no evidence that freight rates are related to physical crude and WTI futures prices differentials.

Another group of studies, primarily from maritime economics, focuses on the physical market for oil. From all the merchandise carried by sea, crude oil accounts for more than 50% in the terms of weight (Stopford, 2009, p. 24). The effect of oil prices in shipping is two-sideways. Firstly, bunker fuel is one of the main cost factors for vessel operators. Secondly, every ship-operator faces the trade-off between speed reductions and as result fuel cost savings and losses because of slow steaming. Ronen (1982) links *oil price* and *vessel speed*, and presents three models, which determine the optimal speed of one engine vessel for a single vessel at a time. From theory, fuel consumption is assumed to be proportional to the third power of its sailing speed (Ryder and Chappell (1979), Álvarez (2009), which was empirically confirmed by Vernimmen et al. (2007)). Vessel speed affects voyage duration, delivery dates and total bunker expenses, and can translate into cost savings for ship operators.

The later work of Ronen (2011) highlights the cost model of increasing fleet size to compensate for speed optimization in the container shipping and still be in line with the service frequency, however, minimize bunker costs. Both Adland and Strandenes (2004) and Jonkeren et al**.** (2012) integrate freight rates in their models. Adland and Strandenes (2004) find that short-run supply elasticity is high when rates are low, since there is marginal capacity available; but low when rates are high and all marginal supply has been exhausted in the VLCC market. Jonkeren et al. (2012) extend Ronen (1982) findings by using data from European inland waterways to estimate the speed elasticity of freight rates and fuel price. Their findings are somewhat consistent with the theory, supporting the fact that freight rates have a positive effect on speed and fuel prices have a negative effect on speed, although the elasticities they measure are far weaker than the theory predicts. Assmann (2012) specialize their study to a particular vessel class. They attempt to measure the relation between bunker prices, freight rates and speed optimization by studying the VCLL movements. Such a relation is assumed to exist from orthodox optimization theory, yet they too find only some support for such a relation, but with elasticity of lower magnitude than they expect.

We will briefly discuss some issues with the model specifications that Assmann and Junkeren used, which we hope to avoid in our specifications. These papers used econometric models based on Ronen's (1982) theoretical optimization model. This gives a function for optimal speed that is then transformed with natural logs, giving models that can be estimated using OLS (Assmann (2012)) or IV (Jonkeren et al. (2012)), which are variations of:

$$\ln traveltime_{it}^{\theta} = \alpha + \beta_1 \ln freigh\, rates_{it} + \beta_2 \ln bunkeroil\ cost_t + u_{it} + \epsilon_{it}$$

Where θ is a coefficient for the efficiency relationship between speed and fuel consumption.

There are at least four possible problems associated with this approach, any of which can explain the lack of strong results.

Firstly, it imposes a myopic static optimization technique to what is clearly a dynamic problem. Ronen assumes that the optimizers consider only the present when they optimize. Nevertheless, if a ship finishes a journey one day earlier, every concurrent journey will also be moved forward in time. Thus, the effect must be calculated as the present value of a finite sum of future payments, which makes this a dynamic problem. Adland et al. (2017) pointed out that the freight market, while highly competitive when viewed in total and in large aggregates of time, can be decomposed into many micro-markets in the short run. In these micro-markets, the lack of immediately available substitutes means that the markets are far from fully competitive. Since ships are scarce resources in such markets, dynamic optimization based on Hoteling's rules, or some equivalent approach, might prove to be a better fit.

Secondly, the model assumes that all optimization is instantaneous in costs and freight rates, so that optimizers at time *t* respond to prices at time *t*. However, the contracts under which ships operate are not instantaneous. Rather, the terms are fixed some time in advance, based on the prevailing market prices at this time. This means that a lagged effect from the variables is expected and might indeed be the only significant effect we expect. A similar argument can be made regarding the bunker oil. While economists realize that the opportunity cost of bunker oil is the price today, actual business agents mostly use budgets as the tool for economic control. Therefore, it is consistent with normal behavior that they would value bunker oil at the price that they paid for it rather than the correct present value. If this is the case, we should also include lagged variables of the bunker oil. Some papers have relaxed this assumption by adding lagged variables, but the optimization paradox persists, so this explanation might be insufficient by itself.

Thirdly, the model makes restrictive assumptions regarding the functional form of the optimizer's problem. There is little empirical proof that vessel speed is correctly modeled using the Cobb-Douglas function, which underpins the regression specification that much of the current research is based on. Rather, an important motivator in choosing the Cobb-Douglas function is its convenient derivation results, which make it easy to use OLS as a regression tool. Moreover, even if the assumptions the model rests on are reasonable, there might still be other factors that influence decision-makers which prevent the relation from holding strictly. This problem is compounded by the potential for endogeneity issues.

Fourthly, the use of OLS and IV estimation methods can lead to endogeneity bias if the explanatory variables are correlated. Their use implicitly assumes that they are all independent. This is highly unlikely in this case. The price for freight and the price of bunker oil are both driven by aggregate macroeconomic forces. If OLS is incapable of separating the effects of the different variables, it will report both as insignificant even if they are both significant.

We believe that the speed optimization issues we have discussed can be resolved and have a strong relevance for our results. We will discuss this further in the Methodology section.

3. Methodology

The goal of this paper is to analyze detailed data for oil shipments to see if there are frictions, so that more detailed data might inform the price of oil. There is no specific theory that models how our data might improve the oil price predictions. To our knowledge, this exact kind of analysis has not been performed before and the lack of theory in this particular case is not surprising. We have therefore used general economic theory to build a theoretical foundation to answer our questions, which we will present later in this section.

The lack of relevant theory means that our approach will be a data mining effort, where we try different specifications. This approach invalidates tests for statistical inference, since fitting the model to the data can give spurious results. To resolve this, we will look for results that are robust to changes in specification. If we find that the same variables are significant in many different specifications, it is unlikely that it is a result of specification bias. In addition to looking for robust results, we will focus our discussion on joint tests for significance. The binary value of either "jointly significant" or "not jointly significant" strongly limits what kind of interpretation of the results we might have, but it is far less sensitive to bias, particularly when it is combined with a discussion of robustness. Our main goal, therefore, is not to make a finished predictor for the oil price, but rather to check if AIS data has robust and jointly significant effect when predicting the oil price.

We also test Kilian's (2009) hypothesis that the model of the oil price should include three factors: expectations of future prices in addition to supply and demand. We triangulate his instrument with one that we construct independently. While we lack a setup to test for causality, we want to see whether our findings are consistent with his model. If they are, it is a strong indication that the precision of oil price estimates can be improved by including controls for expectations. Furthermore, it will add a valuable point to further understanding of the dynamics of the oil market.

Kilian uses the Baltic index and some equivalent measures of his own construct as an instrument for world GDP. He assumes that the freight rates are a direct function of physical transportation and that these are immune from expectations. He further assumes that the residuals from his model must be independent of actual supply and demand, and therefore the result of expectations.

There are several critiques of Kilian's model, some of which implies that there might be significant bias in the results. We therefore address the critique and explain how we will attempt to resolve the possible biases. The first point is that it is impossible to conduct formalized tests to see whether residuals are caused by expectations or by some other factors. Our paper will inform this critique by triangulating his findings; if our results are equivalent, it will act as support to his model. A second point is the fact that shocks to freight rates are transient, much more so than shocks to GDP or the oil price. As such, the increased demand for freight will be met before the increased demand for oil will be, and this will lead to bias in his specification. While this problem will persist also with our specifications, it is not relevant in this thesis, since we do not interpret our residuals as strongly as Kilian does. We only look for robustness of joint tests for significance. Furthermore, we only include this instrument in a VAR model, which explicitly estimates multiple lagged variables of the instrument; therefore, the transience of a shock to vessel opportunity cost is controlled for and will not lead to biases.

We follow Kilian's approach in using ship opportunity cost as an instrument for world production. His reasoning is that demand for ships at any given time is derived from the need to physically move things, and that instantaneous freight rates are therefore close to immune to expectations. We use the opportunity cost for vessels in the same fashion, by measuring vessel speed rather than freight rates, but we relax the assumption that the speed is immune to expectations.

The inclusion of speed allows us to control for expectations. The reason is that we also include traded volumes, the oil price and their interactions, which by the market-clearing property must control for both supply and demand at any given time. The market-clearing property states that, in an efficient market, price will shift to the point where supply equals demand:

$$\begin{cases} D(t) = f(S_t, P_t) \\ S(t) = g(D_t, P_t) \\ D(t) = S(t) \end{cases}$$

Where D is demand, S is supply, P is price and t is a time subscript.

Including information on both Supply/Demand and price is therefore sufficient to control for these factors. If further instruments are significant, that must mean there are other factors than instantaneous supply and demand influencing the market.

We will now discuss the intuition behind this mathematical specification, and why it is reasonable to believe that expectations could have a significant effect on the price. Both the supply and demand for oil show very low short-run elasticities, which means that the price is volatile, which introduces some risk that consumers would prefer to avoid. Oil production requires significant investments of both time and money before production can begin. These investments are required not only in physical oil fields, but in supporting infrastructure, such as skilled labor, oil transport, and oil rigs. These investments are largely specialized to their purpose and therefore sunk once they are made. For example, there are very few alternative uses for an oil platform. The time, resources and the irreversibility of the investments that are required to increase production mean that supply is slow to respond to increasing prices. On the other hand, the magnitude of these investments means that it is rarely efficient to reduce production once the investment has been made, since marginal costs are a small share of the total. In the short-run, a price-taking producer will not reduce production, even if prices decrease considerably. Some producers are not price takers: OPEC dominates global oil production, and some of them act as swing producers in order to maintain revenues, which leads to the negative relation between oil prices and OPEC production rates (Kaufmann et al., 2004).

Demand for oil similarly has a low short-run elasticity. Oil consumers are often barred from substitution of inputs due to sunk investments. For instance, the only way to reduce gasoline consumption from a car is to drive less or to buy a new car, neither of which might be possible in the short run. Compounding this, oil is often an essential factor to economic activity, so that many will have no choice in paying whatever the market demands; if a person depends on their car to get to work, the cost of gasoline would have to be greater than the wages that person would make if the costs were to act as a deterrent to consumption. The combination of low supply and demand elasticities leads to high price volatility, which consumers dislike. This is necessary but not sufficient for expectations to play a large role in the oil price. If consumers could somehow offset the risk, expectations should no longer matter. In order to qualify the inclusion of instruments for expectations in our models, we must show that insurance against oil price risk is costly and impractical, and we will do so now.

The following chart of the oil price for the last ten years illustrates the problem:

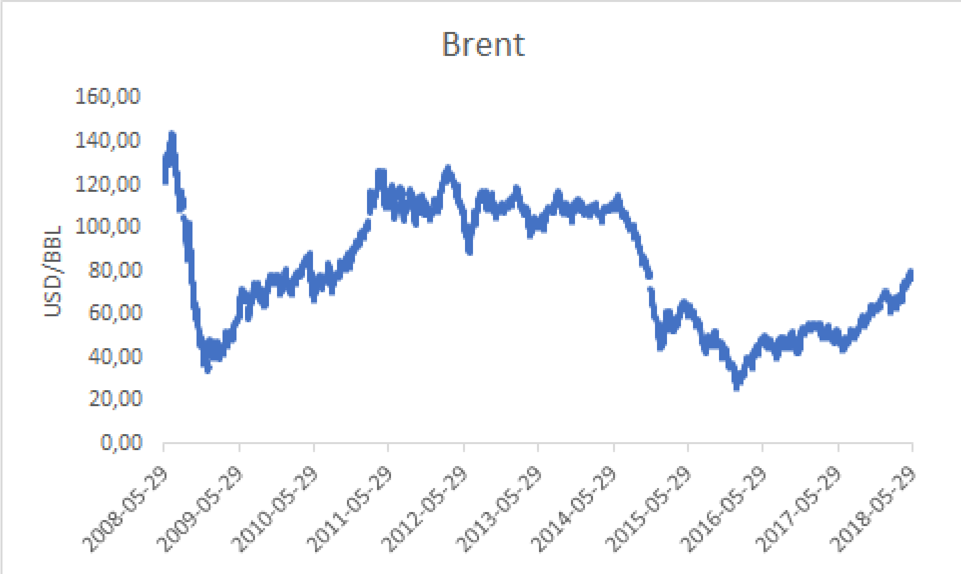

Figure 1: Crude oil prices for Brent (2008-2018)

(U.S. Energy Information Administration, 2018)

As we see from the chart, the oil price has been stable for long periods of time, but with some sudden and wide swings. The risk that a semi-stable price might suddenly change significantly is known as volatility risk and is very expensive to insure against. The reason is that no one has any use for volatility risk; with most derivatives, risk can be sold to someone who wants it to offset some other position. Volatility also complicates the calculation of other risks, as it is often correlated with the kinds of unpredictable shocks that the finance industry abhors. The steep fall in the oil price to the left of the graph, for example, was a result of a depressed outlook caused by the global financial- and debt crisis of '08. This makes volatility expensive to insure against, since the buyer must pay to compensate the seller for the chance of a future shock.

Oil storage could act as an alternative hedge, but it is too costly to be an alternative for anything other than short shocks. The reason is that one would have to store a very large amount of oil to last through a prolonged period of high prices. In figure 1 we see that prices were high for three years in a row, from 2011 to 2014. One would therefore have to store three years' worth of oil in order to hedge against this. This would bind a lot of capital, introduce the risk of spillage and a fire hazard, and require specialized infrastructure. In sum, it would be enormously expensive to store any non-trivial amount of oil for that long, which is why this is not a valid alternative for most consumers.

Since it is difficult to hedge the oil price risk, consumers must attempt to respond to it instead. The market efficiency theory states that any commonly known information must already be priced into the market, so the typical consumer cannot expect to find any more information on the future oil price than the present price. Therefore, the actors closely watch the market, ready to respond if prices start to rise. This generates self-fulfilling prophecies; if the market perceives a positive price signal, demand for futures and derivatives contracts increases as actors try to avoid any further costs; this drives prices up further and prolongs the effect. This is the effect that Kilian describes as expectation-led demand.

3.1 Model 1

A common point-of-departure in econometrics is a basic OLS model. We define the model:

$$Oil\ price_t = \alpha + \beta_i oilflow_{it} + \epsilon_{it}$$

Where *oilflow* is the weekly trade of oil from one region to another. This specification measures the correlation between the interregional flows of oil and the oil price. It is conceptually hard to imagine a causal effect here since the relationship is assumed to be instantaneous. On the other hand, both price and volume transported are decided by common factors; supply, demand and expectations. The F-test for joint significance still has some interest here; it measures whether there is enough and systematic variation in the data that the regional flows inform the oil price. This measure is interesting even if the results cannot be used directly for predictions.

One might expect to find a spurious regression problem here; the oil price is a highly persistent series, and so are some of the *oilflow* series. The potential issue is that the correlation between the two trends is captured by OLS, which gives spuriously strong results. Fortunately, the series are cointegrated in our case, as the residuals from this regression do not exhibit unit root behavior. While spurious co-integration is a possible issue, it is unlikely in our case, since the Dickey-Fuller cointegration test reports non-unit root behavior for all our specifications. It is, therefore, a robust result (please see appendix for the full table of Dickey-Fuller tests).

To relax the assumption of an instantaneous relationship, we include lagged terms of the *oilflow* variable. This is economically sound: since the time span of our data is a week, and it can take several weeks to deliver the oil, we expect that there might be a delayed effect. Furthermore, ships are often contracted some time in advance. A possible issue is the effect of overfitting the data. Since we only have 167 periods. If we were to use 10 lagged variables with our 10 routes we would estimate 101 coefficients, a ratio that almost guarantees overfitting. We therefore limit the number of lagged variables.

3.2 Model 2

We attempt to control for the known information that is the oil price at $t = t - 1$, and include a lagged term for the oil price to the specification:

$$Oil\,price_t = \alpha + \delta Oil\,price_{t-1} + \beta_i oilflow_{it} + \epsilon_{it}$$

A possible problem with this specification is that we have only a few *Oil price* variables to use, and therefore a limited amount of variation in the explained variable. Since the oil price is a highly stationary series, we know that the value at time = t-1 will have a very strong predictive power for the price at time = t. The previous observation of the oil price could therefore crowd out all other variables. This is actually a form of misspecification: OLS assumes that all observations are independent, and the temporal dependence is then captured by the δ coefficient.

3.3 Model 3

It is feasible that the relationship $Oil\,price_t \sim Oil\,price_{t-1}$ can be captured through by its difference. In this case, there might still be some marginal information in the oil volume data. This specification would allow us to avoid the misspecification problem of the previous models. It is possible that the data we have can help in predicting the *change* in the oil price rather than the oil price itself. We therefore measure

$$Oil\,price_t - Oil\,price_{t-1} = \alpha + \beta_i oilflow_{it} + \epsilon_{it}$$

We also vary the amount of lagged values in this specification.

3.4 Model 4

A possible way to avoid the problem of unknown specifications is to use models with weaker restrictions, such as Vector Autoregressive models. Due to computational limits, these models cannot cover too many lags; the ensuing correlation matrices quickly exceed STATA's limits as the number of lagged variables increases. In addition, the vast number of parameters that must be estimated makes overfitting a near certainty if the period is too long. On the other hand, there is no reason to believe that frictions in the markets persist for a long time. In any competitive market, mispricing must eventually adjust itself. The large number of actors in the market, as well as the geographic flexibility that oil shipping involves, should mean that any sub-optimal pricing or oil distribution should resolve itself quickly.

The method descriptions in this part are largely based on Bjørnland & Thorsrud (2015) chapters 7 and 8. The Vector Autoregressive (VAR) model works by fitting a correlation matrix to a set of variables in a way that allows for linear correlation between all variables. For example, a two-variable model with one lag would estimate the relationship

$$\begin{pmatrix} Oil\,price_t \\ Oilflow_t \end{pmatrix} = \begin{pmatrix} \alpha_1 \\ \alpha_2 \end{pmatrix} + \begin{pmatrix} \delta_1 & \beta_1 \\ \beta_2 & \delta_2 \end{pmatrix} \begin{pmatrix} Oil\,price_{t-1} \\ Oilflow_{t-1} \end{pmatrix} + \begin{pmatrix} e_{1,t} \\ e_{2,t} \end{pmatrix}$$

Where β_1 is the interaction between $Oil\,price_t$ and $Oilflow_{t-1}$. This simple two-variable case is equivalent to the OLS model

$$Oil\,price_t = \alpha + \beta\ Oil\,price_{t-1} + \delta\ Oilflow_{t-1} + e_t$$

If we had used more variables, the relationship would not be possible to regress using OLS. The interesting effect is the controls for both the correlation terms β_i and the autoregressive terms δ_i between all variables. This approach removes one of the most awkward assumptions of linear regression: the assumption of linearly independent coefficients. When the true shape of the data is close to linear independence, regression with OLS or equivalent techniques yield good results. However, if the variables in the specification are far from linearly independent, OLS might fail to find the true relationship due to endogeneity. The VAR specification explicitly controls for such interactions.

A possible disadvantage is the loss of an instantaneous relationship,

$$Oilprice_t \sim Oilflow_t.$$

However, since the purpose is prediction of future prices, the excluded information would not have been available. The correlation matrix also implicitly includes the relationship

$$Oilprice_t \sim Oi\widehat{lflow}_t$$

That is the instantaneous relationship with the *predicted* oilflow value.

Economically, there is good reason to expect strong linear dependence of our variables. The world economy is strongly integrated, so a change in demand from South Asia will in most cases impact the demand from Southeast Asia. Furthermore, there is an explicit correlation due to the supply constraints: since supply changes very slowly, increased exports to one area must equal reduced exports somewhere else.

This intuition also extends to the question of ship speed. The assumptions behind ship optimization models explicitly assume that ships respond to bunker oil costs and freight rates. These concepts cannot be independent on a regional basis, due to the fact that both ships and bunker oil can be transferred between shipping routes. Since this kind of substitution happens all the time, the coefficients of OLS will be dependent on each other, and linear decomposition will fail. In a sense, OLS will try to move one part of the metaphorical ship faster than another and find that this cannot be done.

It is worth mentioning that we have chosen not to use a Structured Autoregressive (SVAR) model. Economists often prefer these models because they allow an avenue for economic theory into the model, as well as more interpretable coefficients. If the theoretical assumptions are true, imposing a stricter structure on the data can improve precision, but if the assumptions fail, the results will be biased. The necessary assumption to go from VAR to SVAR is that the lags of some variables cannot influence some other variables, but only be influenced by them. That is, lagged variable A can influence variable B, but not vice-versa. The SVAR estimates a relationship like

$$\begin{pmatrix} Oil\,price_t \\ Oilflow_t \end{pmatrix} = \begin{pmatrix} \alpha_1 \\ \alpha_2 \end{pmatrix} + \begin{pmatrix} \delta_1 & 0 \\ \beta_2 & \delta_2 \end{pmatrix} \begin{pmatrix} Oil\,price_{t-1} \\ Oilflow_{t-1} \end{pmatrix} + \begin{pmatrix} e_{1,t} \\ e_{2,t} \end{pmatrix}$$

In other words, it is assumed that $\beta_1 = 0$. A common approach is to first estimate an unstructured VAR, and then perform the Granger causality test. If this fails to reject the null hypothesis of no Granger causality, one can move forwards with the SVAR specification. In our case, the lack of economic theory regarding the form of our models makes it difficult to defend the use of a SVAR specification. In addition, the Granger tests showed that causality flows in both directions (see appendix for full table).

We estimate the model

$$Oil\ price_t = \alpha + \boldsymbol{\beta oilflow_t} + \epsilon_{it}$$

Bold font is used here to indicate a vector or matrix. Beta loses its subscript as it is now a matrix, and oilflow is a matrix of values as well as lags. We have included up to 4 lagged variables in this specification. In this Vector Auto-Regressive model *beta* is a correlation matrix between the different variables, which are fitted using the Generalized Method of Movements.

3.5 Model 5

We here assess the use of travel time as an instrument using a panel data specification. We have made no logarithmic or other transformations of the data, and because of this, the coefficients are measures of the covariation between the different variables. This correlation is the necessary condition for using vessel speed as an instrument; weak correlation has proved to lead to possible strong biases.

The usual approach is to use ship velocity rather than travel time, but since the shipping lanes have constant length, velocity is simply a scaling of the time spent, since

$$speed = \frac{time}{distance}$$

This transformation makes no difference to the results, except for a scaling of the coefficients.

While economic theory is clear on what the relationship between shipping rates and vessel speed should be, empirical studies so far generally fail to find a significant relationship. This might be a result of the exact specification of the problem; Ronen's (1982) model makes very strong assumptions in order to simplify the optimization. We use the variable in a less constrained setting, so it could still be significant in our case.

We estimate the model

$$Travel\ time_{ti} = \alpha_i + \beta_i cost_{ti} + \lambda bunkeroil_t + \epsilon_{ti}$$

Where t is the time subscript and i is the route subscript. Since we are only interested in whether travel time might be a useful instrument, the question of interest is the significance and impact of the joint β_i coefficients. In order to find comparable freight rates, we had to make some simplifying assumptions. We have assumed that VLCC's are approximate measures of freight costs, which they are as long as the distribution of ships have not changed too much in the period. Equivalently, we have assumed that the distribution of starting- and end points in each region has not changed too much in the period; if they had, the use of travel time in days would be a biased measure for speed. Given the short span of our data, we believe that these assumptions are reasonable.

In the panel data models, we must make a choice between the fixed effects and the random effects estimates. Fixed effects estimate each β_i independently, while random effects use some information from the other variables as well. This increases the available information, which can add to the precision of the results. Also, economic theory assumes that the factors that influence speed velocity are common to all shipping lanes. For these two reasons, it is preferable to use the random effects specification. However, this "pooling" of information makes the random effects specification susceptible to endogeneity, which we suspect might be an issue here. Since both specifications estimate the same model, the coefficients should be equal, if no endogeneity is present. In models 5 and 6 we therefore use the Hausman test, which compares the random effects and fixed effects coefficients, to determine whether random effects is valid or whether we should choose fixed effects instead.

3.6 Model 6

In this model, we replicate the Assmann' s (2012) model:

$$\ln traveltime_{it}^2 = \alpha + \beta_1 \ln freigh\, rates_{it} + \beta_2 \ln bunkeroil\, cost_t + u_{it} + \epsilon_{it}$$

Since

$$\left(\frac{time}{distance}\right)^2 = \frac{time^2}{distance^2}$$

And distance is a constant, squaring travel time is equivalent to squaring velocity, in the sense that tests for significance will be unaffected.

3.7 Model 7

For our final model, we include the speed instrument from model 5 as a measure of the ship's opportunity cost in our full model, in a vector autoregressive model:

$$Oil\ price_t = \alpha + \boldsymbol{\beta oilflowspeed_t} + \epsilon_{it}$$

We exclude the intra-regional oil movements from this specification, since the freight rates for such movements are very uncertain when we aggregate to the regional level. Conceptually, the inclusion of the instrumented speed should act as a control for the opportunity costs of ships. Kilian (2009) uses a similar control based on the Baltic freight rate, where he assumes that the Baltic index is a reasonable instrument for the activity level of the world economy. If his assumption is correct, we should find an equivalent effect, and the inclusion of this instrument might allow us to make better predictions. This model therefore acts as a triangulation of Kilian's results. The oil flows directly measure actual supply and demand, and the opportunity cost of ships, given the actual supply/demand, should capture the expectation effect on oil price.

4. Data

For the purposes of this paper, we used oil tanker shipments data that originates from Clipper Data Ltd. The dataset is a combination of detailed seaborne oil data on individual voyages obtained using Automated Identification System (AIS- tracking), which is used for reporting of real-time vessel position, past-routing and expected port of call; and cargo information from port agents' line-up reports.

Compared to official customs data, which is published monthly, AIS data makes it possible to aggregate information on seaborne oil trade-data at higher frequency. AIS aggregated estimates for seaborne crude exports have proven to be in the alignment with the customs official numbers (Adland et al**.**, 2017). However, at the country level some deviations have been found: countries with storages and transshipment hubs appeared among major exporters, though with no domestic production, and exporters mainly transporting oil by pipeline were not on the list at all. Since we look at the AIS data we get the accurate measures, and the later models 4 and 7 control for interactions. If there is a pattern that is "hidden" by a transshipment hub, the matrix models will therefore find it as interaction between the two flows.

The dataset covers a time-period from January 1, 2013 to mid-March of 2016 (1173 days of observations in total) and contains information on the micro level for 81,728 shipments, including vessel name, vessel class, imo number, loading information (date, load point, port, country, region) and discharge information (date, offtake point, port, country, region). Each observation in the data file starts at the port of loading and ends at the port of discharge. If any of the vessels visits several ports of loading or discharge before arriving at the port of final destination, it is reflected by AIS as individual shipments.

Our data contains an almost excessive amount of information, with details on every trip and every ship in the period. With this level of detail, there are two analytical approaches that can be used. The first is to use all the information to train evolving machine-learning algorithms. The advantage to this method is that it has very few restrictions, and that the format of the data is irrelevant so that no aggregation is necessary. The second approach is to aggregate the data to a form where traditional regression tools are usable. We have chosen the second alternative, for three reasons. First, evolving machine learning is not yet a strong focus at NHH, while the regression techniques are well-established. Second, it is difficult to apply economic theory to an evolving algorithm; while it might yield results that are more precise, those results will not be based on economic theory. Third, the results of evolving algorithms can be difficult or even impossible to interpret.

Since we have chosen the regression approach, we will aggregate the data to a form where regression algorithms can be applied. It is also important to reduce the number of explanatory variables to avoid overfitting and limit the calculation requirements. We have chosen a relatively strong aggregation: we create variables where we sum all oil shipments from one region to another on a week-to-week basis. The weekly aggregation was chosen because we had found that the day-to-day variation of oil shipments was largely noise that contributed nothing to our estimates, and that days with zero shipments were both frequent and caused issues with some of the algorithms. There is an economic argument as well: this variation is caused by the non-divisibility of ships. Oil tankers are very large, and even the largest routes do not require that many tankers each day, so even if daily demand for oil were constant, shipment would vary to use all ship capacity. To further simplify our analysis, we focused on the ten largest inter- and intraregional routes by volume and disregarded the others. These routes represent 54.36 % of volume in the data, which we deemed sufficient.

While we unfortunately lose some of the details that make these data interesting, we saw no other way to approach this problem with traditional techniques. Using country-to-country shipments would lead to 1385 different variables. Since we only have 1173 days of observations in our set, OLS regression techniques would fail to give reasonable results, due to perfectly fitting the data, and multivariate matrix methods would be computationally impossible to solve. Using region-to-region shipments basis reduces this to 178, and it allows us to cover 54 % of the volume by using only ten variables.

Our first use of the vessel speed variable is in our reproduction of Assmann**'s** (2012) model with our data. Since the traveling speed will differ as the ship load changes, we measure the initial trip and return trips as different groups. We then measure how much time each ship spends on each route, and how this changes over time. Since the model of speed optimization expresses optimal speed as an epsilon-root of some expression, where $\epsilon \approx 2$ is a common assumption, we use this to estimate the speed instrument in the regressions for model 6. We therefore square the travel time for this regression. Bunker prices is the 380 CST bunker index, while freight rates are the Clarkson's spot indexes for VLCC's for these routes. Of our ten routes, three do not have costs, since these are intra-regional routes. For the route from the Arab Gulf to South Asia shipments are equally divided between Aframax, Suezmax and VLCC ships. Unfortunately, we only have freight rates for VLCC’s on this route, so that is what we have used. The other routes have 80 % - 90 % VLCC usage, so this should be a good price measure here. We have used the full set of shipments even though not every ship is a VLCC. As long as the share of VLCC’s does not vary too much in the period, along with the oil price, this does not bias the regressions. Given the short time-frame of our data, this is most likely not an issue.

Table 1: Model 1-4. Variables description

Variable	**Obs**	**Mean**	**Std. Dev.**	**Min**	**Max**
Arab Gulf – East Asia	167	9.12	1.22	6.33	12.18
Arab Gulf – North America	167	3.26	1.07	1.44	7.01
Arab Gulf – South Asia	167	3.19	0.87	1.52	5.93
Arab Gulf – Southeast Asia	167	2.39	0.66	0.87	4.89
Eurasia – Northwest Europe	167	1.86	0.39	1.07	3.09
Latin America – Latin America	167	2.30	0.48	1.26	3.72
Latin America – North America	167	1.70	0.40	0.77	3.11
North America – North America	167	1.83	0.45	0.87	3.51
Northwest Europe – Northwest Europe	167	2.27	0.55	0.98	3.91
West Africa – East Asia	167	1.86	0.49	0.76	3.24

Table 2: Model 5-6. Variables description

Variable	Obs	Mean	Std. Dev.	Min	Max
Freight rates	1169	66.71	29.51	18.00	200.00
Bunker oil price	1670	756.93	208.68	314.44	1022.59
Log(freight rates)	1169	4.10	0.46	2.89	5.30
Log(bunker oil)	1670	6.58	0.32	5.75	6.93

Table 3: Model 7. Variables description

Variable	Obs	Mean	Std.dev.	Min	Max
Brent	167	83.13	28.98	27.42	118.05
Trade Volumes					
Arab Gulf – East Asia	167	9.12	1.22	6.33	12.18
Arab Gulf – North America	167	3.26	1.07	1.44	7.01
Arab Gulf – South Asia	167	3.19	0.87	1.52	5.93
Arab Gulf – Southeast Asia	167	2.39	0.66	0.87	4.89
Eurasia – Northwest Europe	167	1.86	0.39	1.07	3.09
Latin America – North America	167	1.70	0.40	0.77	3.11
West Africa – East Asia	167	1.86	0.49	0.76	3.24
Trip Durations					
Arab Gulf – East Asia	167	26.19	2.58	23.63	55.79
Arab Gulf – North America	167	50.07	3.31	42.55	63.31
Arab Gulf – South Asia	167	9.58	3.35	6.56	46.38
Arab Gulf – Southeast Asia	167	19.45	2.39	14.40	28.95
Eurasia – Northwest Europe	167	8.02	1.71	4.08	15.29
Latin America – North America	167	16.93	3.85	10.64	28.24
West Africa – East Asia	167	41.61	3.25	35.87	67.18

5. Results

5.1 Model 1

Table 4: Coefficients for Model 1

Trade route	Coefficient
Arab Gulf – East Asia	-6.776*** (0.000)
Arab Gulf – South Asia	-1.452 (0.530)
Arab Gulf – North America	2.489 (0.177)
Latin America – Latin America	-1.441 (0.729)
North America – North America	-7.133 (0.122)
Northwest Europe – Northwest Europe	-17.47*** (0.000)
Latin America – North America	-6.644 (0.185)
Arab Gulf – Southeast Asia	-4.290 (0.152)
Eurasia – Northwest Europe	1.372 (0.786)
West Africa – East Asia	0.188 (0.963)
Constant	216.3*** (0.000)
N	167
adj. R^2	0.250

p-values in parentheses * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Using this very simple specification, we find that two of the regional flows have highly significant explanatory power for the oil price, and adjusted R^2 is a respectable 25.0%. Since there are two highly significant variables, it is unsurprising that the F-test for joint significance of the *oilflow* variables rejects the null of no significance, with a p-value of 0.

It is important to reiterate that the coefficients in this specification do not measure a casual relation. It is highly unlikely that this specification is the correct one, particularly given the lack of lagged variables and the fact that the mean travel time for the Arab Gulf – East Asia trade route is 26.2, which means that an instantaneous causation would be very surprising. This model basically measures the degree of co-variation between the different trade flows and the oil price, but that is an important result in itself. Multiple significant relationships combined with a non-trivial R^2 is a good starting point for further analysis.

We now extend the analysis by including lagged variables. The literature on lagged variables suggests that an information criterion should be used to find the optimal number of lags. Simply put, we should choose an optimal number of lags, where including one more lag would convey only marginal information. For VAR models it is common to use formal tests such as Akaike's, but for a simple OLS specification, adjusted R^2 is the common measure of model fit, and that is what we use. Based on the following graph, we specify this model with 4 lags.

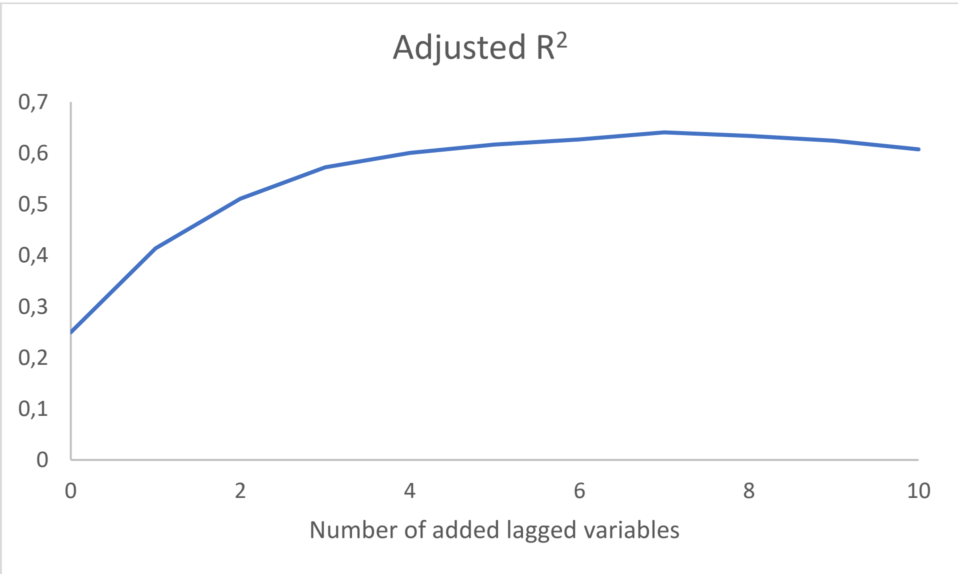

Figure 2: Adjusted R^2 for Model 1

To extend the analysis further, we also regressed 4 lags with the 20 largest routes instead. With an R^2 of 88.0 %, there is every reason to expect overfitting here. To test a hypothesis that some routes and lags are relevant while others are not, we redid the regression using only the significant lags and variables. This should also abate the overfitting problem.

We enclose a subset of the regression tables to illustrate how the changes in specification changes the coefficients. This is the result of redoing the regression using only significant coefficients. Model 1 is the coefficient from the first model, “With 4 lags” is the coefficient from the model with 4 lags, and “Only significant” is a model where only the coefficients that were significant in the 4-lag specification are included.

Table 5: Subset of the regression tables for Model 1

Region	**Lag no°**	**Model 1**	**With 4 lags**	**Only significant**
Arab Gulf - East Asia	0	-6.78***	-4.86***	-6.00***
	1		-5.60***	-7.47***
	2		-6.48***	-8.49***
	3		-6.59***	-8.31***
	4		-5.20***	-6.62***
Arab Gulf - South Asia	0	-1.45	-0.42	
	1		-1.36	
	2		-3.18	
	3		-3.99*	-2.27
	4		-2.91	
Northwest Europe - Northwest Europe	0	-17.47***	-5.32*	-8.35***
	1		-3.13	
	2		-4.43	
	3		-3.49	
	4		-2.85	

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

It is interesting to observe that the coefficients for the Arab-East Asian trade changes so little, even though we removed a large amount of control variables. The Northwest Europe to Northwest Europe instant relationship is similarly consistent. This result is surprisingly robust, in both the significance and the sign; to a certain degree the magnitude of the coefficient as well. As we wrote in the methodology section, we included model 1 and 2 mostly because of their value as measurements for co-variance; we did not expect the coefficients to be unbiased. The consistency of the coefficients of these two regions is therefore both interesting and surprising.

5.2 Model 2

In the first model we deliberately excluded the oil price as an explanatory variable to properly measure co-variation. However, this information is both known and critical to predicting the oil price, and it should therefore be included in any predictive model.

Table 6: Coefficients for Model 2

	Oil price
Lag 1 Oil price	1.000***
	(0.000)
Arab Gulf – East Asia	-0.0125
	(0.935)
Arab Gulf – South Asia	-0.102
	(0.621)
Arab Gulf – North America	-0.209
	(0.204)
Latin America – Latin America	-0.344
	(0.351)
North America – North America	-0.809*
	(0.050)
Northwest Europe– Northwest Europe	0.127
	(0.719)
Latin America – North America	0.121
	(0.786)
Arab Gulf – Southeast Asia	-0.0967
	(0.718)
Eurasia – Northwest Europe	0.613
	(0.173)
West Africa – East Asia	-0.951**
	(0.009)
Constant	3.330
	(0.214)
N	166
adj. R^2	0.994

p-values in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

We find that the explanatory power is extremely high, with R^2 around 99.4 % for both the base model and models including up to 10 lags. We did find that some variables were significant, a few were even consistently significant as the number of lags changed. However, with such a high number of estimated coefficients we should test whether the impact of the variable is significant, not only the statistical tests. As a simple measure, we estimated the model

$$Oil\ price_t = \alpha + \beta\ Oil\ price_{t-1} + \epsilon$$

In other words, we leave out the AIS data entirely.

Table 7: the oil price regressed on its first lag

	Oil Price
Lag 1 Oil Price	1.001***
	(0.000)
Constant	-0.531
	(0.326)
N	166
adj. R^2	0.994

p-values in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The unchanging R^2 shows that the contribution of the AIS data were entirely insignificant in this model. This is in line with our suspicions; since the oil price is autocorrelated, each new observation is not independent of the previous observation. Therefore, we find that the oil price highly correlates with itself, and that correlation dominates the specification entirely.

5.3 Model 3

In this model, we test whether the first differentiated oil price series can be predicted using the AIS data.

Table 8: Coefficients for Model 3

	(1) Δ Brent
Arab Gulf – East Asia	-0.0116 (0.936)
Arab Gulf – South Asia	-0.102 (0.620)
Arab Gulf – North America	-0.209 (0.199)
Latin America – Latin America	-0.344 (0.350)
North America – North America	-0.808^{*} (0.048)
Northwest Europe – Northwest Europe	0.129 (0.695)
Latin America – North America	0.122 (0.782)
Arab Gulf – Southeast Asia	-0.0962 (0.717)
Eurasia – Northwest Europe	0.613 (0.171)
West Africa – East Asia	-0.951^{**} (0.008)
_cons	3.302 (0.134)
N	166
adj. R^2	0.039

p-values in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

We also extend this model using multiple lags. The pattern of adjusted R^2 generated by those models is very interesting:

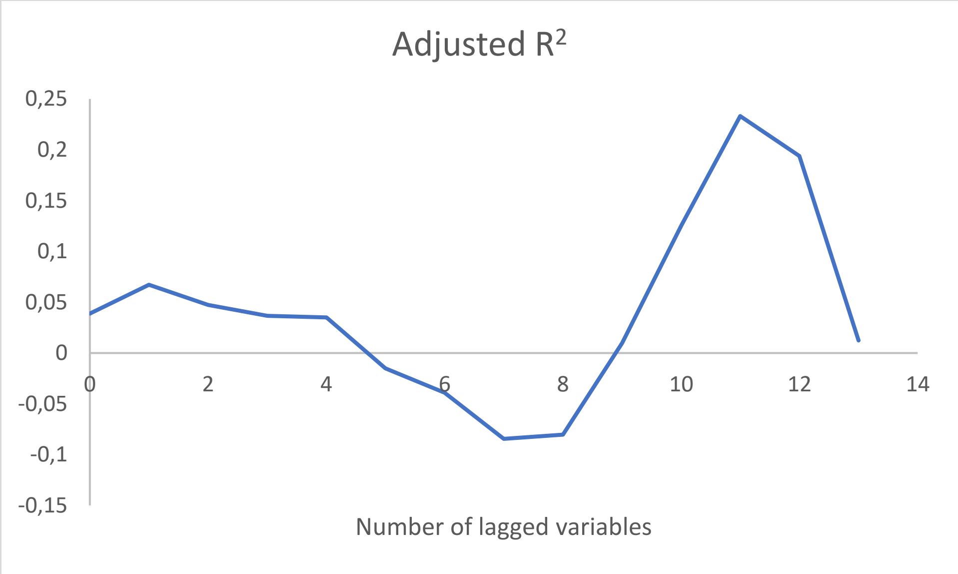

Figure 3: Adjusted R^2 for Model 3

From 14 lags onwards the R^2 explodes, reaching 1 at 15; this is a proof of overfitting; there is possibly an equivalent problem at 10 lags onwards. However, we still find some explanatory power at 0-4 lags; the few variables used here makes overfitting unlikely, as does the plunge from 5 lags to 8. Unfortunately, there are very few significant lags and variables, with a few exceptions that are expected by simple statistics: with 21 estimated coefficients, 2 is expected to be spuriously significant at the 10 % -level, which is what we find. The F-test that all variables have a true value of 0 is 30.99 %, so the model is jointly insignificant.

5.4 Model 4

In a VAR specification it is impossible to estimate an instantaneous model, so we have included the table for four lags instead. The explained variable is Oil price, coefficients and p-values are provided for each explanatory variable in the table.

Table 9: Coefficients for Model 4

Oil Price	**Coef.**	**p-value**	**NW Europe – NW Europe**	**Coef.**	**p-value**
L1.	1.378	0.000***	L1.	-0.144	0.672
L2.	-0.398	0.002***	L2.	-0.473	0.148
L3.	-0.018	0.891	L3.	-0.319	0.321
L4.	0.033	0.684	L4.	0.598	0.075**
Arab Gulf - East Asia			**Latin America – North America**		
L1.	0.019	0.898	L1.	-0.361	0.394
L2.	-0.096	0.533	L2.	0.320	0.445
L3.	-0.081	0.604	L3.	0.067	0.868
L4.	-0.133	0.372	L4.	0.658	0.127
Arab Gulf – South Asia			**Arab Gulf – Southeast Asia**		
L1.	-0.448	0.029**	L1.	-0.194	0.445
L2.	0.120	0.554	L2.	0.123	0.627
L3.	0.160	0.431	L3.	0.097	0.714
L4.	0.213	0.329	L4.	0.513	0.058*
Arab Gulf – North America			**Eurasia – NW Europe**		
L1.	0.324	0.043**	L1.	-0.180	0.673
L2.	-0.036	0.824	L2.	0.400	0.349
L3.	-0.245	0.110	L3.	0.669	0.113
L4.	0.085	0.588	L4.	0.451	0.284
Latin America – Latin America			**West Africa – East Asia**		
L1.	-0.127	0.715	L1.	-0.237	0.480
L2.	0.116	0.729	L2.	-0.501	0.139
L3.	0.122	0.725	L3.	-0.359	0.295
L4.	0.540	0.109	L4.	0.424	0.194
North America – North America			**Constant**	**-2.674**	**0.593**
L1.	0.843	0.047**			
L2.	-0.110	0.780			
L3.	-0.244	0.514			
L4.	-0.247	0.522			

It is of particular interest that only the first two lags of the Oil price itself have statistical significance, and that these have opposite signs. The interpretation of the coefficients here is that a shock to the oil price of USD 1 at t=1 will on average lead to a further increase at t=2, to USD 1.38, with an adjustment downwards to 1 at t=3. This is consistent with the hypothesis that expectations matter; even random noise is initially interpreted as a signal, but the price mostly corrects itself afterwards. It is noteworthy that the "overshooting" of lag 1, which has a value of 1.378-1=0.378 where we subtract the initial shock of 1, is almost equal in magnitude to the correction in lag 2, of 0.398. These results are evidence against a random walk hypothesis, where the coefficient of the oil price lag one should be equal to 1, and every other coefficient should be equal to 0.

There are other variables that have statistically significant coefficients. Of 45 estimated coefficients, we expect around 5 to be spuriously significant at the 5 % level, which is around the number that we find. However, it is interesting that *Arab Gulf – South Asia* and *Northwest Europe – Northwest Europe* both appear to have some statistical significance; they had that in model 2 as well. Equally interesting is that *Arab Gulf – East Asia* has lost all significance here, since it was very consistent in the OLS approach. The question of significance can be further illuminated by the Granger causality results, which we include here:

Table 10: Granger causality results for Model 4

Equation	Excluded Variable	P-value of excluded
Oil Price	Arab Gulf - East Asia	88.0 %
	Arab Gulf – South Asia	22.2 %
	Arab Gulf – North America	19.3 %
	Latin America – Latin America	55.3 %
	North America – North America	32.6 %
	Northwest Europe – Northwest Europe	14.6 %
	Latin America – North America	40.8 %
	Arab Gulf – Southeast Asia	33.2 %
	Eurasia – Northwest Europe	21.5 %
	West Africa – East Asia	24.9 %
	ALL	6.3 %
Arab Gulf – East Asia	Oil Price	0.2 %

As we see, even though every region by itself is insignificant, the joint test of All regions is weakly significant. We also included one line from the table for *Arab Gulf – East Asia*, which explains the apparent paradox of the lost significance. The low p-value indicates that the oil price affects the volume, but not the other way around, at least in this specification. This explains why some OLS specifications picked up a relationship, and also illustrates the value of a more general specification than OLS.

The large number of estimated correlation coefficients makes interpretation somewhat difficult. For example, the third lag of the oil price is not significant, the two previous lags were. Several other variables are significant only in the first lags, but the exact distribution of individually significant lags for a variable might be misleading. It is possible that one data point is sufficient to convey the information from the variable, but exactly which lag will show that information can be extremely unstable; it is possible that only minor changes to the data will change which coefficient will show as the significant one.

Economically, the result is the same. There is no market where the price information from three weeks ago can be viewed isolated from the price two weeks ago. The total interaction of coefficients estimated can be viewed as the way the market responds to shocks, or impulses. There total response is however iterative; a shock to the oil price at t=1 will generate a response to *all other variables* at t=2, which all will generate responses in the oil price in t=3, and so on. The totality of these interactions can be graphed, and those graphs convey more information than the coefficients by themselves. We therefore show the impulse response functions, which is a graphical equivalent of the coefficients in OLS. They show the estimated change to the oil price as a response to a shock to the other variables, with the totality of interactions calculated.

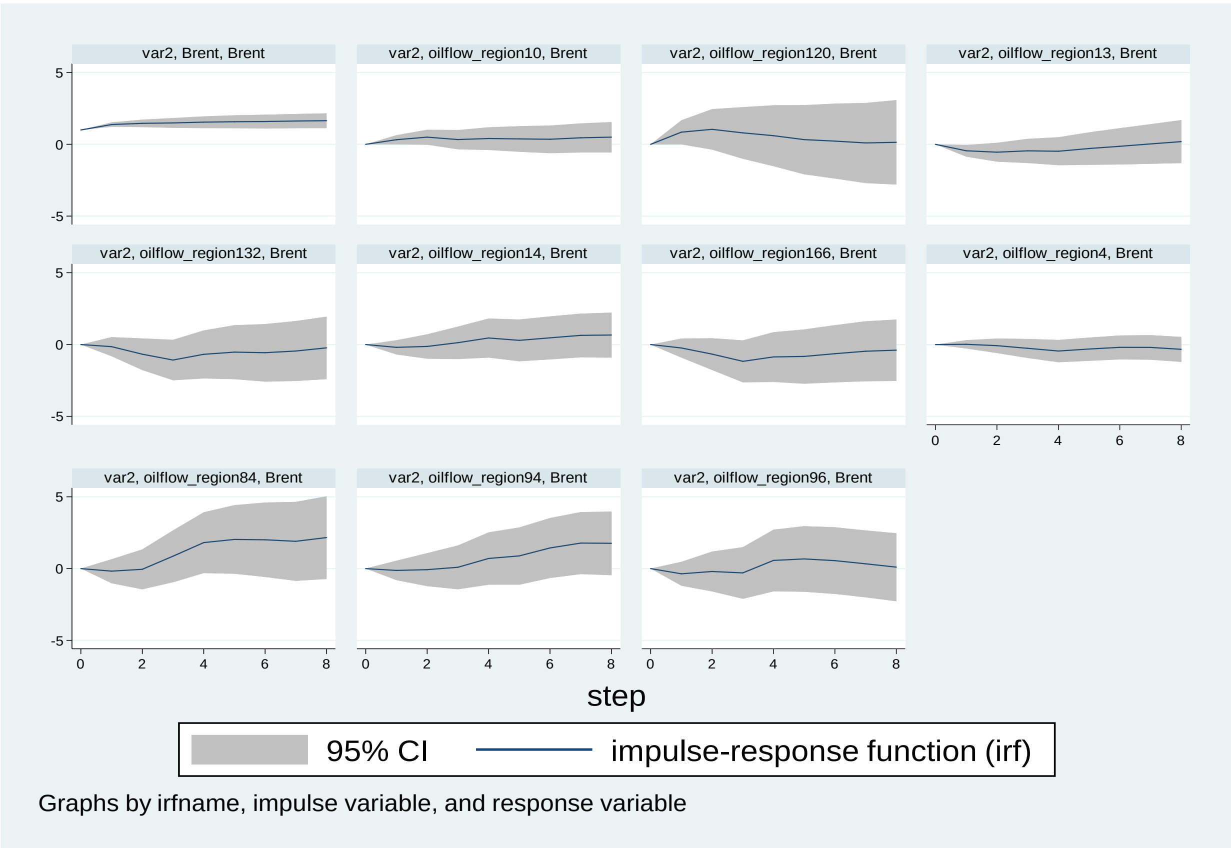

Figure 4: Impulse response functions for Model 4

Table 11: Variables names

Route name on Graph	**Route name from Summary**
flowregion4	Arab Gulf – East Asia
flowregion10	Arab Gulf – North America
flowregion13	Arab Gulf – South Asia
flowregion14	Arab Gulf – Southeast Asia
flowregion84	Eurasia – Northwest Europe
flowregion94	Latin America – Latin America
flowregion96	Latin America – North America
flowregion120	North America – North America
flowregion132	Northwest Europe – Northwest Europe
flowregion166	West Africa – East Asia

As we see from the relatively large movements of the blue lines, the impact from the different regional trade flows is considerable; the lack of explanatory power is a result of the uncertainty of the estimates, expressed by the wide confidence bands. It is noteworthy that most of the blue lines are shaped like either u's or inverted u's. This is consistent with the hypothesis that the market overshoots somewhat after a shock, and then corrects afterwards. Note also that this means that we have temporal endogeneity of the lagged variables: since some of the variation is cancelled out by corrections, even joint tests might fail to pick up a significant relationship, since the totality of the lagged variables goes back to 0. This limits how much joint variation we can measure.

5.5 Model 5

In this model we use a panel data specification to assess the use of travel time as an instrument for the oil price. The Hausman test does not reject the null of no systematic difference in coefficients, so we can choose the less restrictive random effects specifications. We have included a table of 0-4 lags to explore whether our strategy of including lags is viable.

Table 12: Coefficients for Model 5

Included lags	(0) Trip time	(1) Trip time	(2) Trip time	(3) Trip time	(4) Trip time
Freight rates	-0.00677 (0.166)	-0.00836 (0.327)	-0.00950 (0.269)	-0.00991 (0.250)	-0.0118 (0.199)
Bunker price	-0.000609 (0.194)	-0.00638 (0.295)	0.00193 (0.781)	0.000873 (0.901)	-0.00176 (0.815)
L.freight_rates		0.00124 (0.886)	0.00754 (0.515)	0.00794 (0.494)	0.00697 (0.574)
L2.freight_rates			-0.00671 (0.442)	-0.00283 (0.808)	-0.00202 (0.871)
L3.freight_rates				-0.00563 (0.521)	-0.00166 (0.894)
L4.freight_rates					-0.00854 (0.361)
L.bunkerprice		0.00578 (0.344)	-0.0191 (0.110)	-0.0121 (0.341)	-0.00663 (0.628)
L2.bunkerprice			0.0166^{*} (0.017)	-0.00220 (0.863)	-0.0134 (0.354)
L3.bunkerprice				0.0129 (0.068)	0.0346^{*} (0.011)
L4.bunkerprice					-0.0138 (0.068)
Constant	25.46^{***} (0.000)	25.45^{***} (0.000)	25.47^{***} (0.000)	25.55^{***} (0.000)	26.37^{***} (0.000)
N	1169	1162	1155	1148	1141

p-values in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 13: Tests of joint significance for Model 5

Lags	**p-value for joint significance of freight rates**	**Overall R^2**
0	17.86 %	21.91 %
1	32.41 %	19.25 %
2	38.56 %	14.13 %
3	42.80 %	14.23 %
4	18.27 %	18.71 %
5	0.00 %	29.21 %
6	0.00 %	30.22 %
7	0.00 %	30.84 %
8	0.00 %	31.47 %
9	0.00 %	32.00 %
10	0.00 %	32.44 %

With this specification, we find what many others in the literature have already found, a weak to non-existing relationship between freight rates and vessel speed. We find very few significantly significant variables, and no robustness to different specifications. The table of tests for joint significance also confirms that the strategy of including lags is insufficient to solve this problem. The sudden drop at 4 to 5 included lags is combined with a large jump in overall R^2, which is a clear indication of overfitting. This specification thus fails to give any significant results.

5.6 Model 6

In this model we replicate Assmann (2012) model, transforming all variables using the logarithmic function. We keep our panel data setup. The Hausman test strongly rejects the null of no systematic difference in coefficients, so we must choose the fixed-effect specification. We include a table of 0-4 here as well, with the same reason as before.

Table 14: Coefficients for Model 6

Number of lags	(0) Log (travel time)	(1) Log (travel time)	(2) Log (travel time)	(3) Log (travel time)	(4) Log (travel time)
Log(freight)	-0.116** (0.002)	-0.0810 (0.259)	-0.0851 (0.236)	-0.0851 (0.235)	-0.0875 (0.223)
Log(bunker)	-0.100** (0.002)	0.0255 (0.944)	0.497 (0.226)	0.447 (0.280)	0.237 (0.570)
L.Log(freight)		-0.0431 (0.552)	0.0652 (0.505)	0.0757 (0.437)	0.0735 (0.450)
L2.Log(freight)			-0.125 (0.091)	-0.0721 (0.463)	-0.0599 (0.541)
L3.Log(freight)				-0.0791 (0.282)	-0.0319 (0.744)
L4.Log(freight)					-0.0667 (0.363)
L.Log(bunker)		-0.135 (0.710)	-1.508* (0.032)	-1.262 (0.096)	-0.751 (0.330)
L2.Log(bunker)			0.902* (0.030)	0.273 (0.719)	-0.621 (0.451)
L3.Log(bunker)				0.424 (0.318)	2.071** (0.008)
L4.Log(bunker)					-1.077* (0.012)
Constant	7.122*** (0.000)	7.217*** (0.000)	7.297*** (0.000)	7.407*** (0.000)	7.623*** (0.000)
N	1169	1162	1155	1148	1141

p-values in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 15: Tests of joint significance for Model 6

Lags	**p-value for joint significance of travel time**	**Overall R^2**
0	0.22 %	23 .49 %
1	0.72 %	23 .82 %
2	0.48 %	23 .38 %
3	0.47 %	23 .89 %
4	0.69 %	22 .94 %
5	1.01 %	21 .21 %
6	0.18 %	23 .18 %
7	0.18 %	23 .69 %
8	0.30 %	23 .14 %
9	0.46 %	23 .67 %
10	0 .69 %	24 .20 %

The result of this model, like the previous one, is in line with the literature on this subject. We find some significant variables, but their estimated coefficients are far below the estimates we would expect from the optimization theory. Models 5 and 6 jointly show that our strategy of including lagged variables was insufficient to resolve the speed optimization paradox.

5.7 Model 7

Table 16: Coefficients for Model 7

Oil Price	Coef.	p-value	Constant Coef.	p-value
L1.	1.356	0.079*	16.081	0.092*
L2.	-0.452	0.127		
L3.	0.008	0.126		
L4.	0.075	0.082*		
Oil Volumes			**Trip Duration**	
Arab Gulf - East Asia				
L1.	-0.001	0.138	-0.236	0.116
L2.	-0.181	0.149	-0.110	0.124
L3.	-0.105	0.157	-0.233	0.124
L4.	-0.154	0.150	-0.060	0.118
Arab Gulf - North America				
L1.	0.311	0.150	0.092	0.055*
L2.	0.030	0.156	-0.112	0.051*
L3.	-0.210	0.145	-0.006	0.053*
L4.	0.188	0.151	-0.022	0.053*
Arab Gulf - South Asia				
L1.	-0.477	0.205	0.110	0.090*
L2.	0.047	0.202	0.063	0.093*
L3.	-0.214	0.209	0.113	0.093*
L4.	-0.128	0.205	0.102	0.094*
Arab Gulf - Southeast Asia				
L1.	0.069	0.241	0.070	0.068*
L2.	0.220	0.236	0.154	0.068*
L3.	0.175	0.236	-0.133	0.067*
L4.	0.315	0.239	-0.012	0.068*
Eurasia - Northwest Europe				
L1.	0.546	0.434	-0.064	0.107
L2.	0.239	0.449	0.096	0.110
L3.	0.843	0.411	0.039	0.111
L4.	0.481	0.424	0.146	0.102
Northwest Europe – Northwest Europe				
L1.	-0.881	0.455	-0.029	0.045**
L2.	0.429	0.443	0.069	0.047**
L3.	-0.169	0.416	-0.012	0.048**
L4.	1.006	0.417	-0.051	0.050*
West Africa - East Asia				
L1.	-0.241	0.325	-0.047	0.051*
L2.	-0.791	0.346	-0.008	0.052*
L3.	-0.727	0.352	-0.073	0.052*
L4.	0.105	0.320	0.112	0.051*

In this model, we combine the speed instrument with a VAR model. We find that 5 of 7 vessel speed measures are statistically significant. It is interesting that there is so little variation of p-values, across all variables. No trade flows are found to be significant. However, the Granger results in table 17 show a different picture. The left p-value is the value of whether the variable Granger predicts the oil price, the right p-value is the value of whether the variable is Granger predicted by the oil price.

When we moved on to the VAR from the OLS models, we were worried that endogeneity might be preventing OLS from finding the correct results. We pointed out that both the supply of oil and the supply of ships is limited, which means that a change in one regional variable must be correlated with other regional variables. Also, the demand for oil in different regions is driven by the same economic forces and will thus be correlated; the same can be said about shipping speeds. The Granger test from model 7 show that we were right to worry.

Table 17: Coefficients for Model 7

Variable	**p (-> Brent)**	**p (<- Brent)**
Oil volumes		
Arab Gulf - East Asia	0.645	0.000*
Arab Gulf - North America	0.092*	0.387
Arab Gulf - South Asia	0.061*	0.442
Arab Gulf - Southeast Asia	0.635	0.186
Eurasia - Northwest Europe	0.068*	0.030**
Northwest Europe – Northwest Europe	0.016**	0.023**
West Africa - East Asia	0.070*	0.015**
Trip Durations		
Arab Gulf- East Asia	0.031**	0.999
Arab Gulf - North America	0.085*	0.000*
Arab Gulf - South Asia	0.118	0.770
Arab Gulf - Southeast Asia	0.029**	0.019**
Eurasia - Northwest Europe	0.444	0.186
Northwest Europe – Northwest Europe	0.459	0.000**
West Africa - East Asia	0.138	0.003**

In this specification, 5 of 7 trade flows statistically predict the oil price. The flow from the Arab Gulf to East Asia does not predict the price but responds instead. This is consistent with the results from model 4, and it is interesting that this result is so robust. Only one trade flow neither predicts or is predicted by the oil price. The same holds for the speed instrument; 3 of 7 instruments predict the oil price, and 5 of 7 responds to it; only one is neither predicted or predictive.

The large number of significant variables is a highly interesting result, particularly since the VAR specification also contains all the information from the lagged oil price observations. Moreover, if we consider the following chart of Impulse Response Functions, it is clear that many of the trade routes have significant impacts as well as statistical significance:

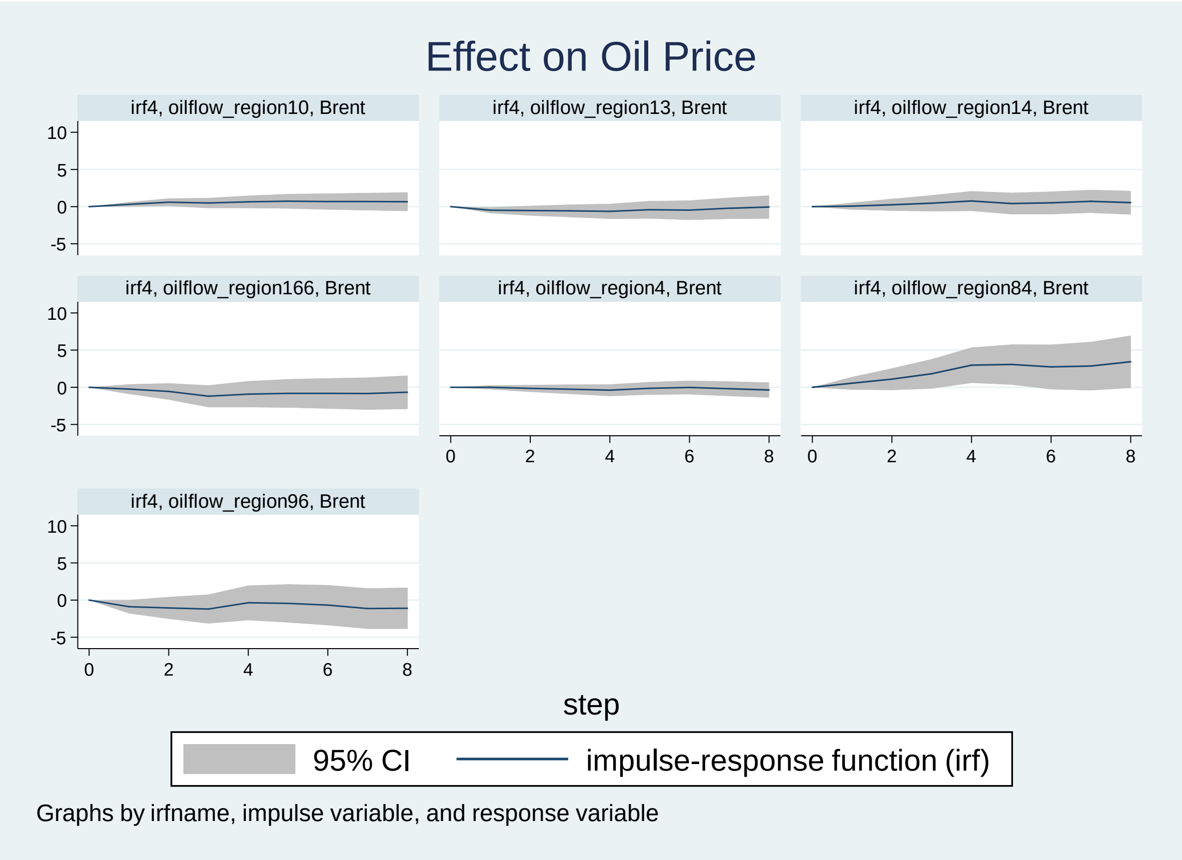

Figure 5: Impulse response functions for Model 7

Equivalent results hold for the speed instrument, which is included in the appendix. Also in the appendix is the full table of the Granger causality tests. It is far too long to include here, but it is very interesting. We find a large degree of endogeneity in the speed instrument; 3 variables are significantly influenced by multiple other speed variables. This can help to explain the lack of significance in models 5 and 6, that persisted even as we added lags; in this model we control for lags and endogeneity, and significance increases dramatically. This clearly points to endogeneity as a cause of the lack of significance. We discuss comparisons between models 4 and 7 in the Discussion section.

We also find that the speed instrument both affects and is affected by the oil price. The two-way speed-oil price causality that we found is consistent with Kilian's (2008) model, since an underlying assumption there is that shipping rates can be used as an instrument for global economic activity. The instrument on speed that we have used here is another measure of the opportunity cost of the ships. The fact that the causality between instrumented speed and oil price goes both ways is indicative of some underlying factor that affect both; the activity of the world economy seems a likely suspect. While we cannot address the unbiasedness of Kilian's instrument, our findings are at least consistent with his assumptions, and thus act as a triangulation of same.

6. Discussion and concluding remarks

We have explored the relationship between the flow of crude oil and the oil price. Tangentially, we have also examined the relationship between ship speed and freight rates. We have also offered a triangulation of Kilian's (2009) assumption that opportunity costs of a ship can be used as an instrument for world GDP.

It is common to use measures of model fit to asses which model performs better, such as Root Square Mean Error, but we have chosen not to do so. Since our dataset covers only a few years, and the oil price is influenced by macroeconomic trends that stretch longer than that, we are incapable of assessing the precision of our models without potentially creating large biases in the measures, since even a well-crafted model might find results that are spurious to the period. This means that it is impossible to assess the validity of our models out-of-sample. If longer data is available sometime in the future, it might be worth re-investigating.

We have found that it is feasible that information from the AIS tracking system be useful in predicting the oil price, if used in a proper specification. OLS models failed to give any significant results when information from the oil price itself was included, which could indicate that there was no more information to extract from the AIS data. However, we have argued that this is probably a result of misspecification bias; both the assumptions of independent observations and endogeneity were violated, which could suppress otherwise significant results. VAR models gave mixed results. In model 4 we failed to find a strong relationship, but we did find a weak link. The Granger p-value of 6.3 % is not as strong as we would have liked to have for a conclusive result, but there exists some statistical relation. Model 7, which included the vessel speed instrument, gave much stronger results.

We have also examined the apparent paradox of non-optimal ship speed behavior. We offered some critique of the current models and tried a specification that should solve the issued we raised. We found that a less stringent specification apparently resolved the paradox; the opportunity cost of ships and the oil price are consistently correlated when we control for both lags and endogeneity. The short time-span is again preventing us from conclusively saying that the issue is resolved, but it appears at least to be worthy of further investigation.

In the same model, we assessed the validity of using ship opportunity cost as a measure for GDP. While we cannot address the question of possible bias created by Kilian's use of the

Baltic index, we nonetheless offer conceptual support, as our unrelated instrument for the same opportunity cost showed strong statistical significance.

The most interesting result is found by comparing models 4 and 7; in model 7 we use vessel speed as an instrument to control for the role of expectations; in model 4 we do not. Unfortunately, we cannot isolate the effect of this instrument. The VAR specification is too unrestricted to allow interpretation of some coefficients in isolation, which is one reason why many economists prefer SVAR specifications. We have chosen not to use the SVAR models, since we suspected that endogeneity was a problem, which could create bias in a SVAR model. The high degree of two-way Granger causality between the different variables is proof of endogeneity. The fact that most variables had high Granger causality significance and low coefficient significance is supportive evidence. This result throws doubt on the unbiasedness of Kilian's results, as they rely on the SVAR specification.

Despite this, our findings support the overlying concept of Kilian's model. While we cannot measure the precise effect of the speed instrument, we can compare the models with and without it. Model 7 stands out from our other models in the number of significant variables and the strength of the Granger causality tests. A majority of the Granger tests in model 7 are significant, both among the oil trade and ship speed variables. This is consistent with Kilian's model. Since our results are constructed in an entirely different and independent matter, this triangulation is a supportive result of his general model.

We can draw other conclusions from comparing the VAR with and without the speed instrument. First, the market is not perfectly identified using only Quantity and Price, because if it were, the vessel speed instrument would not be significant. As a continuation of this argument, the market for oil is not perfect, because if it were, Quantity and Price would be sufficient information to identify the market. While these are not the most exciting conclusions, it is nevertheless good to see that the prices are more complex than that, as it validates our attempt to find deeper relations. We can further conclude that inclusion of AIS, and by extension other types of non-standard data, can be a viable approach. Most of these observations were highly significant in the last VAR model, which also included all data from the oil price; it follows that they must have added information.

The significance of the speed instrument also supports the theory of vessel speed optimization. Although we did not test this relationship directly, we indirectly used vessel speed as a measure of the opportunity cost of ships. The fact that this information was significant shows that there must exist some relationship between vessel speed and opportunity cost, which freight rates also measure. The lack of significance in models 5 and 6 indicates that controlling with lagged variables was insufficient to resolve the paradox. The increase in significance to model 7, combined with the strong indications of endogeneity, is a strong indication that endogeneity lies at the heart of this paradox. As far as we know, the use of a vessel speed instrument directly in a VAR model has not been done before in this setting, so it is possible that we have found a genuine discovery.

We must, however, be careful not read too much into the results; they are interesting, but not proof of the actual effect. The short time-span of our dataset makes these predictions dubious. It is possible that the two effects that we just discussed is a mere artifact from that period; the oil price increased, and exports from the Arab Gulf shifted from North America to South Asia. We meet here a constant problem in macroeconomic modeling: we only have one draw of outcomes from the world. The macro shift towards Asia is the only data that we will ever have, so it is impossible to find counterfactual observations. As such, even if the model fits very well, it might be the result of overfitting. Usually, we would not expect overfitting to be a problem with only 4 weeks of lagged variables. However, if the macro trends are strong enough, they might lead to overfitting even with very sparse specifications.

For future avenues of research, the endogeneity issues that we found in the vessel speed instrument stands out. This can possibly explain the apparent vessel speed optimization paradox. A further exploration of this issue is an exciting possibility. Another option is expanding the scope of time for this kind of study. The results would be more reliable if they covered a larger amount of time. This would also allow for reasonable estimations of model fit, which we lack in our results. Possibly, this could lead to more precise estimates of the oil price, which could have large economic consequences.

7. References

Adland, R., & Strandenes, S. P. (2004). A discrete-time stochastic partial equilibrium model of the spot freight market. *Discussion Paper 11/2004, Norwegian School of Economics (NHH)*.

Adland, R., Haiying, J., and Strandenes, S. P. (2017). Are AIS-based trade volume estimates reliable? The case of crude oil exports. *Maritime Policy & Management, 44 (5)*, 657-665.

Alizadeh, Amir H.; Nomikos, Nikos K. (2004). Cost of carry, causality and arbitrage between oil futures and tanker freight markets. *Transportation Research Part E, Vol.40(4)*, 297-316.

Álvarez, J. F. (2009). Joint Routing and Deployment of a Fleet of Container Vessels. *Maritime Economics & Logistics, Vol.11(2)*, 186.

Alvarez-Ramirez, Jose; Alvarez, Jesus; Rodriguez, Eduardo. (2008). Short-term predictability of crude oil markets: A detrended fluctuation analysis approach. *Energy Economics, Vol.30(5)*, 2645-2656.

Alvarez-Ramirez, Jose; Cisneros, Myriam; Ibarra-Valdez, Carlos; Soriano, Angel. (2002). Multifractal Hurst analysis of crude oil prices. *Physica A: Statistical Mechanics and its Applications, Vol.313(3)*, 651-670.

Assmann, L. M. (2012). *Vessel speeds in response to freight rate and bunker price movements : an analysis of the VLCC tanker market.* Bergen: Norwegian School of Economics (NHH).

Bjørnland, H., & Thorsrud, L. (2015). *Applied time series for macroeconomics.* Oslo: Gyldendal akademisk.

BP. (2018). *Energy Outlook* .

Charles, Amélie; Darné, Olivier. (2009). The efficiency of the crude oil markets: Evidence from variance ratio tests. *Energy Policy, Vol.37(11)*, 4267-4272.

Clarksons Research. (u.d.). Hentet fra https://www.clarksons.net/

Clipper Data Ltd. (2016). *www.clipperdata.com*.

Coleman, L. P. (2012). Explaining crude oil prices using fundamental measures. *Energy Policy, Vol.40*, 318-324.

Dées, Stéphane; Karadeloglou, Pavlos; Kaufmann, Robert K.; Sánchez, Marcelo. (2007). Modelling the world oil market: Assessment of a quarterly econometric model. *Energy Policy, Vol.35(1)*, 178-191.

Elder, John; Serletis, Apostolos. (2008). Long memory in energy futures prices. *Review of Financial Economics, Vol.17(2)*, 146-155.

Ewing, Bradley T.; Harter, Cynthia Lay. (2000). Co-movements of Alaska North Slope and UK Brent crude oil prices. *Applied Economics Letters, Vol.7(8)*, 553-558.

Ewing, Bradley T.; Thompson, Mark A. (2007). Dynamic cyclical comovements of oil prices with industrial production, consumer prices, unemployment, and stock prices. *Energy Policy, Vol.35(11),* 5535-5540.

Green Steven L. and Mork Knut Anton. (1991). Toward Efficiency in the Crude-Oil Market. *Journal of Applied Econometrics, 6(1)*, 45-66.

Jonkeren, Olaf; Ommeren, Jos Van; Rietveld, Piet. (2012). *Freight prices, fuel prices, and speed.(Report).* Journal of Transport Economics and Policy, Vol.46(2), p.175(14).

Kaufmann, Robert K.; Dees, Stephane; Gasteuil, Audrey; Mann, Michael. (2008). Oil prices: the role of refinery utilization, futures markets and non-linearities. *Energy Economics, Vol.30(5*, 2609-2622.

Kaufmann, Robert K.; Dees, Stephane; Karadeloglou, Pavlos; Sanchez ; Marcelo. (2004). Does OPEC matter? An econometric analysis of oil prices. *The Energy Journal, Vol.25(4)*, 67(24).

Kaufmann, Robert K.; Ullman, Ben. (2009). Oil prices, speculation, and fundamentals: Interpreting causal relations among spot and futures prices. *Energy Economics, Vol.31(4)*, 550-558.

Kilian, L. (2009). Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks. *American Economic Review, Vol. 99 Issue 3*, 1053-1069.

Lin, Sharon Xiaowen; Tamvakis, Michael N. (2001). Spillover effects in energy futures markets. *Energy Economics, Vol.23(1)*, 43-56.

Maslyuk, Svetlana; Smyth, Russell. (2008). Unit root properties of crude oil spot and futures prices. *Energy Policy, Vol.36(7)*, 2591-2600.

Mensi, Walid; Aloui, Chaker; Hamdi, Manel; Nguyen, Duc Khuong. (2012). Crude oil market efficiency: An empirical investigation via the Shannon entropy. *International Economics, Vol.129,* 119(19).

Milonas, Nikolaos T.; Henker, Thomas. (2001). Price spread and convenience yield behaviour in the international oil market. *Applied Financial Economics, Vol.11(1),* 23-36.

Ronen, D. (1982). The Effect of Oil Price on the Optimal Speed of Ships. *The Journal of the Operational Research Society, Vol.33, No. 11*, 1035-1040.

Ronen, D. (2011). The effect of oil price on containership speed and fleet size. *The Journal of the Operational Research Society, Vol. 62, No.1*, 211-216.

Ryder, S.C., Chappell, D. (1979). Optimal Speed and Ship Size for Liner Trades. *Marine Transport Centre, University Of Liverpool*.

Serletis, Apostolos; Andreadis, Ioannis. (2004). Random fractal structures in North American energy markets. *Energy Economics, Vol.26(3),* 389-399.

Silvapulle Param and Moosa Imad A. (1999). The relationship between spot and futures prices: Evidence from the crude oil market. *The Journal of Futures Markets, Vol. 19, No. 2,* 175-193.

Stopford, M. (2009). *Maritime economics.* London: Routledge.

Tabak, Benjamin M.; Cajueiro, Daniel O. (2007). Are the crude oil markets becoming weakly efficient over time? A test for time-varying long-range dependence in prices and volatility. *Energy Economics, Vol.29(1),* 28-36.

U.S. Energy Information Administration. (2018, June 19). *Crude Oil Prices: Brent - Europe [DCOILBRENTEU], retrieved from FRED, Federal Reserve Bank of St. Louis*. Hentet fra https://fred.stlouisfed.org/series/DCOILBRENTEU

Vernimmen, B., Dullaert, W., & Engelen, S. (2007). Schedule Unreliability in Liner Shipping: Origins and Consequences for the Hinterland Supply Chain. *Maritime Economics & Logistics*, *Vol.9(3)*, 193.

8. Appendix

Table 18: Cointegration tests: Augmented Dickey-Fuller on residuals

Level	1%	5%	10%			
Critcal values	-3.489	-2.886	-2.576			
Lags	Model 1	Model 2	Model 2b	Model 3	Model 4	Model 7
0	-3.843	9.441	-8.779	-9.440		
1	-4.256	-9.479		-9.458		
2	-4.273	-9.738		-9.746		
3	-3.805	-9.572		-9.620		
4	-3.221	-9.370		-9.389	-12.081	-12.461
5	-3.026			-9.702		
6	-3.035			-9.772		
7	-3.243			-9.625		
8	-2.917			-10.203		
9	-2.951			-10.424		
10	-3.603			-9.900		

Models 5 and 6 - Dickey Fuller not possible for panel data

Table 19: Granger causality tests from model 4

Equation	Excluded	p-value
Brent	oilflow_region4	0.88
	oilflow_region13	0.222
	oilflow_region10	0.193
	oilflow_region94	0.553
	oilflow_region120	0.326
	oilflow_region132	0.146
	oilflow_region96	0.408
	oilflow_region14	0.332
	oilflow_region84	0.215
	oilflow_region166	0.249
	ALL	0.063*
oilflow_region4	Brent	0.002***
	oilflow_region13	0.552
	oilflow_region10	0.396
	oilflow_region94	0.102
	oilflow_region120	0.326
	oilflow_region132	0.04**
	oilflow_region96	0.809
	oilflow_region14	0.329
	oilflow_region84	0.531
	oilflow_region166	0.158

	ALL	0,000
oilflow_region13	Brent	0.914
	oilflow_region4	0.14
	oilflow_region10	0.043**
	oilflow_region94	0.7
	oilflow_region120	0.966
	oilflow_region132	0.319
	oilflow_region96	0.424
	oilflow_region14	0.214
	oilflow_region84	0.38
	oilflow_region166	0.438
	ALL	0.111
oilflow_region10	Brent	0.764
	oilflow_region4	0.619
	oilflow_region13	0.1
	oilflow_region94	0.105
	oilflow_region120	0.396
	oilflow_region132	0.527
	oilflow_region96	0.103
	oilflow_region14	0.51
	oilflow_region84	0.331
	oilflow_region166	0.439
	ALL	0.047**
oilflow_region94	Brent	0.923
	oilflow_region4	0.58
	oilflow_region13	0.024**
	oilflow_region10	0.038**
	oilflow_region120	0.694
	oilflow_region132	0.155
	oilflow_region96	0.329
	oilflow_region14	0.652
	oilflow_region84	0.388
	oilflow_region166	0.919
	ALL	0.137
oilflow_region120	Brent	0.015**
	oilflow_region4	0.045**
	oilflow_region13	0.001***
	oilflow_region10	0.004***
	oilflow_region94	0.369
	oilflow_region132	0.543
	oilflow_region96	0.079**
	oilflow_region14	0.001***

	oilflow_region84	0.324
	oilflow_region166	0.499
	ALL	0.000***
oilflow_region132	Brent	0.113
	oilflow_region4	0.129
	oilflow_region13	0.038**
	oilflow_region10	0.914
	oilflow_region94	0.981
	oilflow_region120	0.009***
	oilflow_region96	0.092*
	oilflow_region14	0.036
	oilflow_region84	0.198
	oilflow_region166	0.052*
	ALL	0.000***
oilflow_region96	Brent	0.192
	oilflow_region4	0.008***
	oilflow_region13	0.006***
	oilflow_region10	0.018**
	oilflow_region94	0.94
	oilflow_region120	0.254
	oilflow_region132	0.591
	oilflow_region14	0.038**
	oilflow_region84	0.497
	oilflow_region166	0.316
	ALL	0.004***
oilflow_region14	Brent	0.59
	oilflow_region4	0.426
	oilflow_region13	0.759
	oilflow_region10	0.557
	oilflow_region94	0.572
	oilflow_region120	0.157
	oilflow_region132	0.186
	oilflow_region96	0.567
	oilflow_region84	0.402
	oilflow_region166	0.969
	ALL	0.266
oilflow_region84	Brent	0.035**
	oilflow_region4	0.106
	oilflow_region13	0.00***
	oilflow_region10	0.468
	oilflow_region94	0.479
	oilflow_region120	0.723

	oilflow_region132	0.223
	oilflow_region96	0.119
	oilflow_region14	0.081*
	oilflow_region166	0.349
	ALL	0.004**
oilflow_region166	Brent	0.024**
	oilflow_region4	0.477
	oilflow_region13	0.991
	oilflow_region10	0.021**
	oilflow_region94	0.006***
	oilflow_region120	0.045**
	oilflow_region132	0.154
	oilflow_region96	0.601
	oilflow_region14	0.002***
	oilflow_region84	0.094*
	ALL	0.002***

Figure 6: IRF for the time instrument for model 7

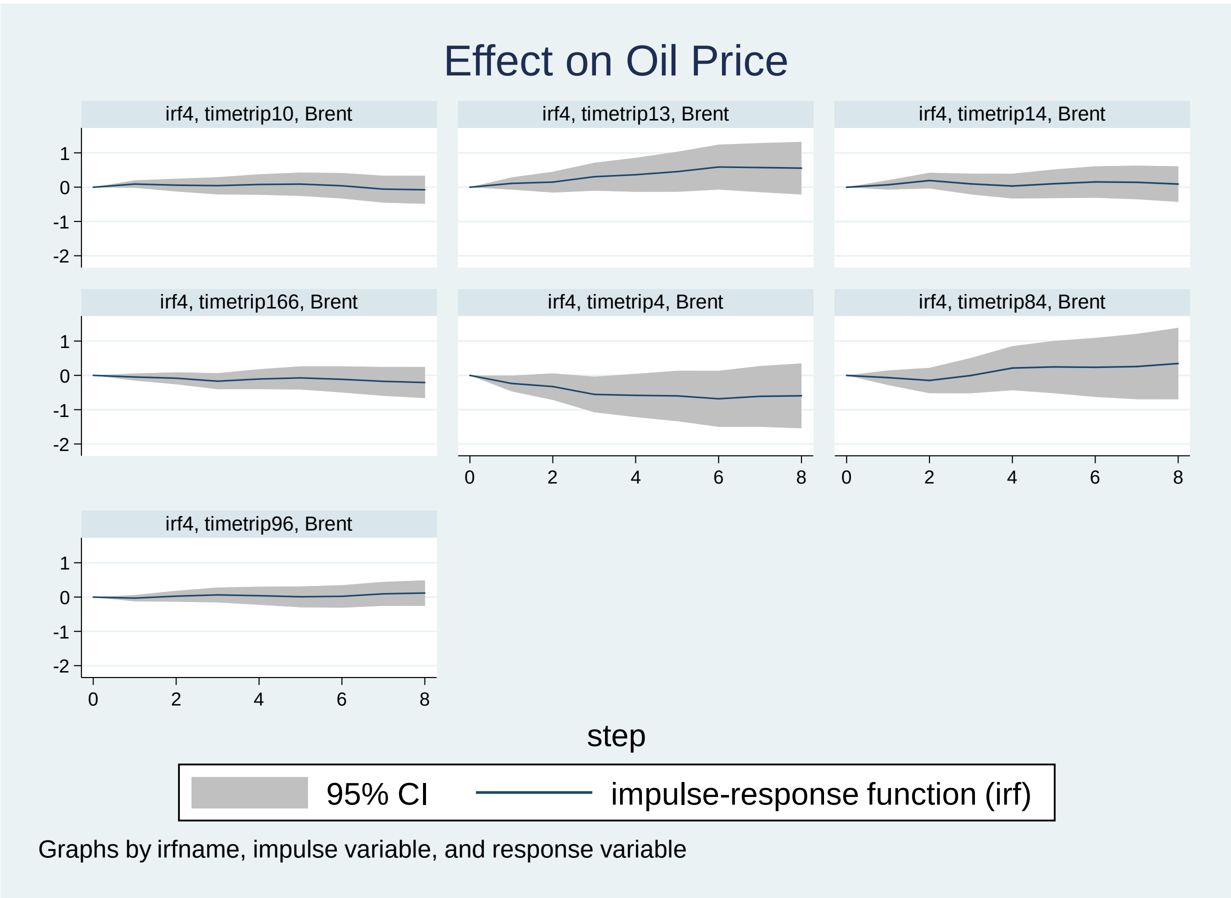

Table 20: Granger causality tests for model 7

Equation	Excluded	p-value
Brent	oilflow_region4	0.645
	oilflow_region10	0.092*
	oilflow_region13	0.061*
	oilflow_region14	0.635
	oilflow_region84	0.068*
	oilflow_region96	0.016**
	oilflow_region166	0.07*
	timetrip4	0.031**
	timetrip10	0.085*
	timetrip13	0.118
	timetrip14	0.029**
	timetrip84	0.444
	timetrip96	0.459
	timetrip166	0.138
	ALL	0.000***
oilflow_region4	Brent	0.000***
	oilflow_region10	0.026**
	oilflow_region13	0.306
	oilflow_region14	0.584
	oilflow_region84	0.323
	oilflow_region96	0.32
	oilflow_region166	0.05*
	timetrip4	0.562
	timetrip10	0.067
	timetrip13	0.548
	timetrip14	0.003
	timetrip84	0.216
	timetrip96	0.18
	timetrip166	0.856
	ALL	0.000***
oilflow_region10	Brent	0.387
	oilflow_region4	0.786
	oilflow_region13	0.193
	oilflow_region14	0.792
	oilflow_region84	0.242
	oilflow_region96	0.574
	oilflow_region166	0.433
	timetrip4	0.173
	timetrip10	0.659
	timetrip13	0.696

	timetrip14	0.173
	timetrip84	0.38
	timetrip96	0.753
	timetrip166	0.393
	ALL	0.027**
oilflow_region13	Brent	0.442
	oilflow_region4	0.001***
	oilflow_region10	0.002***
	oilflow_region14	0.557
	oilflow_region84	0.174
	oilflow_region96	0.104
	oilflow_region166	0.147
	timetrip4	0.148
	timetrip10	0.013**
	timetrip13	0.075*
	timetrip14	0.59
	timetrip84	0.002***
	timetrip96	0.003***
	timetrip166	0.07*
	ALL	0.000***
oilflow_region14	Brent	0.186
	oilflow_region4	0.01*
	oilflow_region10	0.253
	oilflow_region13	0.61
	oilflow_region84	0.412
	oilflow_region96	0.109
	oilflow_region166	0.384
	timetrip4	0.171
	timetrip10	0.076*
	timetrip13	0.197
	timetrip14	0.344
	timetrip84	0.096
	timetrip96	0.091*
	timetrip166	0.207
	ALL	0.005***
oilflow_region84	Brent	0.03**
	oilflow_region4	0.207
	oilflow_region10	0.18
	oilflow_region13	0.000***
	oilflow_region14	0.003***
	oilflow_region96	0.029
	oilflow_region166	0.054
	timetrip4	0.189

	timetrip10	0.015**
	timetrip13	0.04**
	timetrip14	0.309
	timetrip84	0.128
	timetrip96	0.324
	timetrip166	0.263
	ALL	0.000***
oilflow_region96	Brent	0.023**
	oilflow_region4	0.001***
	oilflow_region10	0.017**
	oilflow_region13	0.001***
	oilflow_region14	0.034**
	oilflow_region84	0.047
	oilflow_region166	0.019**
	timetrip4	0.000***
	timetrip10	0.002**
	timetrip13	0.002**
	timetrip14	0.344
	timetrip84	0.283
	timetrip96	0.482
	timetrip166	0.055*
	ALL	0.000***
oilflow_region166	Brent	0.015**
	oilflow_region4	0.362
	oilflow_region10	0.011**
	oilflow_region13	0.942
	oilflow_region14	0.031**
	oilflow_region84	0.047**
	oilflow_region96	0.408
	timetrip4	0.899
	timetrip10	0.168
	timetrip13	0.623
	timetrip14	0.672
	timetrip84	0.068*
	timetrip96	0.797
	timetrip166	0.538
	ALL	0.103
timetrip4	Brent	0.999
	oilflow_region4	0.434
	oilflow_region10	0.076*
	oilflow_region13	0.098*
	oilflow_region14	0.775
	oilflow_region84	0.706

	oilflow_region96	0.612
	oilflow_region166	0.105
	timetrip10	0.619
	timetrip13	0.13
	timetrip14	0.103
	timetrip84	0.504
	timetrip96	0.942
	timetrip166	0.601
	ALL	0.176
timetrip10	Brent	0.000***
	oilflow_region4	0.053*
	oilflow_region10	0.000***
	oilflow_region13	0.016**
	oilflow_region14	0.006***
	oilflow_region84	0.002***
	oilflow_region96	0.081*
	oilflow_region166	0.493
	timetrip4	0.015**
	timetrip13	0.002***
	timetrip14	0.081
	timetrip84	0.318
	timetrip96	0.055*
	timetrip166	0.049**
	ALL	0.000***
timetrip13	Brent	0.77
	oilflow_region4	0.491
	oilflow_region10	0.062
	oilflow_region13	0.083*
	oilflow_region14	0.898
	oilflow_region84	0.733
	oilflow_region96	0.508
	oilflow_region166	0.076*
	timetrip4	0.139
	timetrip10	0.161
	timetrip14	0.205
	timetrip84	0.34
	timetrip96	0.253
	timetrip166	0.465
	ALL	0.083*
timetrip14	Brent	0.019*
	oilflow_region4	0.001***
	oilflow_region10	0.523
	oilflow_region13	0.000***

	oilflow_region14	0.016**
	oilflow_region84	0.000***
	oilflow_region96	0.01**
	oilflow_region166	0.17
	timetrip4	0.023**
	timetrip10	0.015**
	timetrip13	0.865
	timetrip84	0.607
	timetrip96	0.000***
	timetrip166	0.517
	ALL	0.000***
timetrip84	Brent	0.44
	oilflow_region4	0.000***
	oilflow_region10	0.752
	oilflow_region13	0.737
	oilflow_region14	0.032**
	oilflow_region84	0.291
	oilflow_region96	0.162
	oilflow_region166	0.166
	timetrip4	0.321
	timetrip10	0.841
	timetrip13	0.508
	timetrip14	0.135
	timetrip96	0.183
	timetrip166	0.147
	ALL	0.000***
timetrip96	Brent	0.000***
	oilflow_region4	0.283
	oilflow_region10	0.069*
	oilflow_region13	0.026**
	oilflow_region14	0.391
	oilflow_region84	0.025**
	oilflow_region96	0.854
	oilflow_region166	0.076*
	timetrip4	0.362
	timetrip10	0.623
	timetrip13	0.737
	timetrip14	0.49
	timetrip84	0.403
	timetrip166	0.941
	ALL	0.000***
timetrip166	Brent	0.003**
	oilflow_region4	0.081*

oilflow_region10	0.003***
oilflow_region13	0.011**
oilflow_region14	0.65
oilflow_region84	0.003***
oilflow_region96	0.17
oilflow_region166	0.539
timetrip4	0.114
timetrip10	0.004
timetrip13	0.091*
timetrip14	0.114
timetrip84	0.531
timetrip96	0.245
ALL	0.000***